Naveen Jindal School of Management

2014-04

A Psychological Reexamination of the Bertrand Paradox

UTD AUTHOR(S): Ernan E. Haruvy

Southern Economic Journal 2014, 80(4), 948–967
DOI: 10.4284/0038-4038-2012.264

A Psychological Reexamination of the Bertrand Paradox

Enrique Fatas,* Ernan Haruvy,† and Antonio J. Morales‡

The Bertrand paradox describes a situation in which two competing firms reach an outcome where both price at marginal cost. In laboratory experiments, this equilibrium is not generally observed. Existing empirical works on Bertrand competition have found evidence for boundedly rational models. We find that such models are useful in organizing behavior in early stages of the game, but less so in later stages. We show that a new model, coarse grid Nash equilibrium, based on the assumption that subjects discretize the strategy space, explains the data better.

JEL Classification: C91, D43, L13

1. Introduction

The Bertrand Paradox describes a situation in which two firms engaging in price competition reach an outcome where both charge a price equal to the marginal cost. This is because under the Bertrand setting each firm has a strong incentive to undercut the price of its opponent by a small increment to capture the entire market. The outcome is called a paradox because the two firms could easily earn positive profits by charging a higher price. While charging a higher price than marginal cost would not be a pure strategy equilibrium for strictly rational selfish players, it is an intuitive outcome for such settings, and one that is consistent with the empirical studies of duopoly competition we discuss below.

In an attempt to reconcile theory and empirical evidence, economists have followed different approaches. On one hand, it has been shown that the Bertrand paradox is not generally robust to minor modifications to the Bertrand setting. Realistic features include capacity constraints (Edgeworth 1925), quantity pre-commitments (Kreps and Scheinkman 1983), differentiation (Singh and Vives 1984), uncertainty (Janssen and Rasmusen 2002), and nonconstant marginal costs (Dastidar 1995; Abbink and Brandts 2008; Argenton and Müller 2012). Any of these modifications result in both positive earnings and an interior solution. Thus, one way to obtain more reasonable outcomes is to modify the theoretical Bertrand setting.

---

* School of Economics, University of East Anglia, Norwich Research Park, Norwich, NR4 7TJ, United Kingdom; E-mail e.fatas@uea.ac.uk.

† School of Management, University of Texas at Dallas, 500 W Campbell Road, SM32, Richardson, TX 75080-3021, USA; E-mail eharuvy@utdallas.edu.

‡ Department of Economics, University of Malaga, Facultad de Ciencias Economicas, Departamento de Teoria Economica, Plaza El Ejido s/n, 29013 Malaga, Spain; E-mail amorales@uma.es.

We acknowledge the support of Spanish MEC grants ECO2010-20584 and ECO2011-26996 and the ESRC Network for Integrated Behavioral Science.

Received November 2012; accepted October 2013.

Another possibility is the relaxation of the assumptions regarding the rationality of the economic agents to arrive at empirically supported results. This is where laboratory experiments are very useful. The laboratory has long been a useful test bed for theories pertaining to oligopoly competition (e.g., Chamberlin 1948; Fouraker and Siegel 1963; Apesteguia, Huck, and Oechssler 2007; Abbink and Brandts 2008; Apesteguia et al. 2010; Argenton and Müller 2012). In the laboratory we can have greater control of the underlying properties of the competitive setting, so that we can focus on behavioral explanations. Indeed, in the laboratory, when we keep the strict Bertrand assumptions, we still do not observe the Bertrand equilibrium. This article places itself in this latter avenue. However, rather than assuming that firms depart from the Nash equilibrium assumption, as is common in behavioral theories offered to explain these observed patterns, we propose that firms select prices on a coarser grid.

Why would firms be willing to use coarser rather than finer grids? One possible explanation is a cognitive limitation explanation, driven by the cognitive difficulty involved in making decisions on the continuum. An active debate in cognitive science concerns the issue of whether perception is continuous or discrete. As VanRullen and Koch (2003) state, "it seems surprising that such a fundamental question … has not been definitively answered one way or another." From a cognitive perspective, the discreteness of perception might be a feature of human thinking and therefore should be considered as a source of bounded rationality. This human tendency, demonstrated when reading scales (Mitchell 2001), can in some cases be attributed to a conscious human decision to simplify. Computations with round numbers are easier to perform and help alleviate the informational requirements of complex computations.[1] This explanation (e.g., Rubinstein 1998) has received strong psychological support (Payne et al. 1988) and can also be considered as an example of categorization, identified by Gigerenzer et al. (1999) as a building block of bounded rationality, and recently applied in game theoretic settings by Jehiel[2] (2005).

How can we demonstrate that the coarse grid Nash equilibrium explanation is more appropriate for explaining the Bertrand paradox than other approaches? In a linear Bertrand setting, the coarse grid explanation adds two additional Nash equilibria above marginal cost. These additional equilibria imply two additional local modes at particular distances from the marginal costs, and they are not predicted or supported by other behavioral explanations.

We run an experiment with three treatments and two different Bertrand games. Both games are played repeatedly and share a strategy space with admissible prices in the interval [0, 500], large enough for subjects to have the opportunity to use coarse grids. In both games we consider a duopoly with fixed demand; the unique difference is the cost function: In one case we analyze constant marginal costs, and in the second we consider quadratic costs. This second game is designed to avoid the zero profit condition in equilibrium, which has been highlighted as a negative feature of the basic linear Bertrand setting. Both games are played using a fixed matching protocol, as the composition of markets does not change across rounds. As an additional control, we also run a last variant of the constant marginal costs version of the Bertrand game. In this last treatment markets are reshuffled at the end of each round, as a way to control for any repeated game effect.

[1] In line with this, Harris (1991) argues that traders may coordinate to restrict the price set to a limited discrete set of prices to limit negotiations and save deliberation resources as well as reducing the probability of costly mistakes.

[2] The equilibrium concept in Jehiel (2005) applies to multistage games with perfect information, rather than simultaneous move games of the kind studied here, although the Bertrand game is repeated.

Table 1. Summary of Experimental Design

	Linear Bertrand		Quadratic Bertrand
Matching	Fixed	Random	Fixed
Subjects	46	90	84
Sessions	2	2	2
No. independent observations	23	9	42
Rounds	20		20
Price range	0–500		0–500
Profit for lowest price in pair	$20(P-5)$		$20P-5\times20^2$
Profit for highest price in pair	0		0
Profit in case of tie	$10(P-5)$		$10P-5\times10^2$
Feedback after each round	Both prices and profits		Both prices and profits

We find that in both games subjects seem to use coarse grid as the price distribution in both games is characterized by spikes situated at convenient (from a human perspective) round prices: 50, 100, 150, or 200. Despite these similarities, the outcomes in the two games are quite different. We find that the quadratic Bertrand game—the more complicated problem of the two—results in rapid convergence to the efficient equilibrium. This is in marked contrast to both versions of the linear Bertrand game, where no convergence to equilibrium is observed, even when prices decline slightly faster in the treatment with random rematching.

In seeking an explanation for why we observe these differences between the games, we test our proposed coarse grid Nash equilibrium together with two popular boundedly rational approaches: Quantal Response Equilibrium (Baye and Morgan 2004) and step-level thinking (Gneezy 2005). We estimate the three models and compare them in terms of their likelihood. In the linear version of the Bertrand game as well as a quadratic cost variation, we find that a coarse grid Nash equilibrium is a better predictor of prices, relative to competing models, in each of the three treatments. This is good news as the coarse grid Nash equilibrium naturally explains multimodal price distributions.

The rest of the article is organized as follows. Section 2 contains the experimental design and procedures. The behavioral theories are presented in section 3. Section 4 presents the experimental results as they relate to the behavioral conjectures. Section 5 concludes.

2. Experimental Design and Procedures

Our experiment consists of two games and three different treatments. Table 1 contains a summary of them.

The Games

We consider two Bertrand duopoly games with a homogeneous product, no capacity constraints, inelastic demand $Q>0$, and complete information. The market is served by the lowest price firm or split in the case of a tie. The two games differ on the cost function. In the linear cost game, the cost function for quantity q is

$$c_i(q)=cq \text{ for } i=1,2, \tag{1}$$

whereas in the quadratic case the cost function is

$$c_i(q) = cq^2 \text{ for } i = 1,2. \tag{2}$$

As a useful benchmark, we briefly characterize the standard Nash equilibria in the two Bertrand duopoly games we study, with prices in the interval $[0, \infty)$.

PROPOSITION 1. (Proof for quadratic case is in the Appendix.) The set of Nash equilibria is

i. For the linear Bertrand setting, $p^* = c$
ii. For the quadratic game, $p^* = \left[\frac{1}{2}cQ, \frac{3}{2}cQ\right]$.

For the linear Bertrand setting, the standard prediction of "price is equal to marginal cost" holds, and equilibrium profits are zero. The quadratic case is different, as there is a multiplicity of Pareto-ranked equilibria. The smallest equilibrium price of $\frac{1}{2}cQ$ is the lowest in the Pareto ranking of equilibria and yields zero profit, whereas the largest equilibrium price of $\frac{3}{2}cQ$ yields the efficient equilibrium.

Experimental Design

The basic laboratory version of the Bertrand model, the linear Bertrand game, has fixed demand and constant costs. The linear Bertrand game has been shown to deviate from Nash equilibrium in the laboratory (Dufwenberg and Gneezy 2000, hereafter DG). A key feature of the linear Bertrand game is that in equilibrium competitors earn zero or near zero profits and thus is often not reached in practice. DG note that this feature is "an undesirable feature, which one might have suspected would undermine the attraction of the equilibrium outcome" (DG, p. 19, footnote 9). This explains why in addition to the linear Bertrand setting, we include a quadratic Bertrand game with quadratic costs that does not have this undesirable feature.

Abbink and Brandts (2008) and Argenton and Müller (2012) proposed Bertrand games with quadratic marginal costs and linear demand. Such settings involve a range of equilibria and therefore present an interesting problem of equilibrium selection. Although these settings do indeed shed light on how subjects coordinate, the incentives to collude on the joint profit-maximizing price (the cartel outcome) are strong and the cartel outcome is commonly observed. In contrast, we eliminate the linear demand feature of their setting, opting instead to restore the fixed demand feature of the linear Bertrand game. This makes the quadratic Bertrand game we use more comparable to the linear Bertrand game and serves to make the collusive outcome far less salient and sustainable.[3]

We choose $Q = 20$ and $c = 5$ as a compromise between the linear and the quadratic games. Because the marginal cost is positive but small, we get close to the standard setting with null marginal cost in the linear game (as in DG), while we allow for positive equilibrium prices in the quadratic game (the interval [50,150]). The set of admissible prices is the interval [0, 500], so the

[3] Argenton and Mueller (2012, p. 511) explain that in their setting the riskiness of equilibrium increases in the size of demand. This is because the potential loss from miscoordination increases with the size of demand. High Bertrand prices associated with collusion are safer in their setting because the unilateral benefits from undercutting are relatively small. This argument does not apply to fixed demand Bertrand settings.

unique standard Nash equilibrium price for the linear game $p^* = 5$ is quite close to zero and represents only one-hundredth of the length of the price interval.

The experiment was framed as price competition. To ensure the appeal of discretization as a simplification device, subjects were not informed about the parameter values or the precise demand functions—only about the general features of the demand (that was fixed) and cost (linear or increasing, depending on the treatment).

In the experiment, subjects played the duopoly game for 20 rounds. In the first treatment subjects face an environment equivalent to the one used by DG (with linear costs and fixed matching). The other two treatments can be viewed as different controls. As explained above, the quadratic Bertrand experiment has no corner equilibrium solutions. The linear Bertrand experiment with random matching helps us to discern for the existence of a repeated game effect. In this last treatment, subjects were randomly grouped in cohorts of 10 players at the beginning of the experiment. They were randomly rematched with a different player from their cohort at the end of each round. Even when the number of rounds exceeds the cohort size, the probability of being rematched with the same player in any posterior round is still very low.

As the results of both variants of the linear Bertrand game, with fixed or random matching, were very similar, the quadratic duopoly was run only under a fixed matching scheme. In all three treatments, subjects were informed at the end of each round of their opponent's posted prices and profits. A table containing past information was also available.

Experimental Procedures

Experiments were run in the laboratory for research in experimental economics at the University of Valencia (LINEEX). A total of 220 students from business and economics were recruited using a standard electronic recruitment procedure. Forty-six students played the linear Bertrand game with fixed matching, 90 students participated in the linear Bertrand game with random matching, and the remaining 84 played the quadratic game with fixed matching. Two sessions were run for each treatment, and subjects earned an average pay of $25 for approximately one hour of participation.

3. Behavioral Predictions

Before presenting the experimental results, we review in this section the predictions of the proposed coarse grid Nash equilibrium together with two popular behavioral models that have been proposed to account for the Bertrand paradox: Quantal response equilibrium and step-level thinking.

Extant Behavioral Predictions

The Quantal Response Equilibrium

Quantal response equilibrium (QRE) is based on the notion that subjects make errors associated with the expected payoff of each alternative, and these errors follow a Type I extreme value distribution. The result is a probabilistic decision rule, which maps expected payoffs to choices via a logit. Expected payoffs are derived based on players' beliefs about the strategies of other players. The equilibrium notion means that expectations about the choices of competitors likewise follow a logit structure. Baye and Morgan (2004) applied QRE to DG's Bertrand competition data and have shown it to perform well.

Denote expected payoff to a price of x by Π_x and denote the observed empirical probability of event E by $P(E)$. We compute expected payoffs for each of the 501 possible prices based on the observed empirical frequency in the game (for the quadratic case we have):

$$E\Pi_x = (20x - 2000)P(price > x) + (10k - 500)P(price = x). \tag{3}$$

We then compute a logit best response function as

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

where $\hat{P}(price = x)$ is the estimated probability of observing price x. We substitute the computed logit probabilities from the last stage for the empirical probabilities. We recompute the logit probabilities and repeat the procedure until 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.

Step-Level Thinking

Step-level thinking is a popular approach to bounded rationality, based on the seminal article of Stahl and Wilson (1994). Recently Fatas and Morales (2013) show that these models of strategic thinking are able to explain the behavior in costly coordination games. The starting point of that theory is a level-1 player who holds a naive uniform prior—believing that all feasible actions by his opponent are equally likely. The theory also allows for player types with more sophisticated reasoning abilities that follow a hierarchical structure. For example, a level-2 type anticipates others to be of level-1, a level-3 might anticipate that others follow level-2, and so on.

Gneezy (2005) reported a Bertrand price competition experiment with 10 price points from 1 to 10. Using the step-level thinking framework with some modifications, he found that the model could explain Bertrand paradox behavior with remarkable fit. In contrast to Gneezy (2005), we have a range of 501 possible choices. In that range, the difference between level-1 and level-2 (and higher) is too small to be meaningful. Specifically, the likelihood does not budge significantly when adding level-2 and higher. 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.

The Coarse Grid Nash Equilibrium (CGNE)

A prediction of prices equal to marginal costs is a consequence of the assumption that decision makers adjust prices using the minimal feasible decrement. This assumption, extensively used as epsilon-type arguments, does not seem to describe human behavior in all situations, as distinctive price patterns around convenient round numbers are found in a variety of environments.

A well-known phenomenon in the marketing of consumer goods is *psychological pricing*, that is, prices just below some round number. As a consequence, there is a majority of prices that have a 9 or a 5 as the last significant digit (see Friedman 1967 for a classic reference and Schindler and Kirby 1997). A related phenomenon observed in stock markets is price clustering: stock prices cluster on whole numbers (Osborne 1965 and Niederhoffer 1965 and 1966 are classic references). A number of recent articles provides further confirmation across markets: Harris (1991) on the NYSE, Ascioglu, Comerton-Forde, and Mclnish (2007) on the Tokyo stock exchange, and Sonnemans (2006) on the Dutch stock market. In the willingness-to-pay literature, there is evidence of a prominence effect: When people are asked to provide an estimate using a payment scale, they tend to select prominent values on that scale (Whynes, Phillips, and Frew 2005).

The patterns of discretization discussed above are also evident in auctions. The auction literature is especially relevant because the simple Bertrand setting we are interested in is akin to a simple two-player auction. Coarse grid bidding is commonly observed in auction settings. This pattern of behavior is evident in both the experimental lab (Neugebauer and Selten 2006)[4] and the field. Under strict assumptions of a simplified auction environment, price bids in increments larger than the minimum feasible may appear dominated. However, since such "jump bids" (as they have sometimes been termed in the literature) are so commonly observed, theories have emerged to attempt to explain them as a rational decision making outcome. Isaac, Salmon, and Zillante (2007) document jump bidding in data from FCC 3G spectrum license auctions. They construct a model in which jump bidding occurs due to strategic concerns as well as impatience. Kwasnica and Katok (2007) showed that when bidders are impatient—the opportunity cost of time is higher in one experimental manipulation in that study—bidders respond by choosing larger jump bids.

The coarse grid model assumes that subjects simplify the strategy space and consider only prices that are multiples of k; for example, the original strategy space is replaced by the set $\{0,k,2k,3k\ldots\}$. It can be shown that the coarse grid adds two additional equilibrium prices above the standard NE. Proposition 2 refers to the linear cost game (proof in the Appendix).

PROPOSITION 2. Consider the linear cost model and a common grid $k>0$. The set of coarse grid Nash equilibria is (p, p) with $p=nk$ for n integer such that $\frac{c}{k}\leq n\leq 2+\frac{c}{k}$.

The conditions $\frac{c}{k}\leq n\leq 2+\frac{c}{k}$ imply that there will be at least two multiples of k that are CGNE of the linear model (because the difference between the upper and the lower bounds is precisely 2). And in certain cases, there would be a third CGNE that will coincide with the NE of the linear game. This happens whenever the ratio c/k is a non-negative integer: 0, 1, 2 … . Examples of this situation are (i) when the marginal cost is zero and, therefore, $c/k=0$, and the three CGNE are 0, k, and $2k$ and (ii) when k is a divisor of the marginal cost, and the corresponding three CGNE are c, $c+k$, and $c+2k$. This latter case generalizes the known result that with prices restricted to be integers ($k=1$), there are equilibria at c, $c+1$, and $c+2$.[5] The novelty of our approach is that the discretization is not externally imposed but is rather a mental grid of the strategy space.

In the quadratic cost case, we can state the analogous proposition (proof in the Appendix) regarding equilibria for a common grid of size k. Define $n^{-}(\cdot)$ and $n^{+}(\cdot)$ as follows:

[4] They report that 58% of subjects' first bids are in multiples of 5, for example, exactness of 5, and that this fact could be explained by prominent number considerations.

[5] See, for example, footnote 1 in DG.

$$n^{-}(k) = min\{n|nk \geq 50\}, \tag{5}$$

$$n^{+}(k) = max\{n|nk \leq 150\}. \tag{6}$$

Then we can state proposition 3.

PROPOSITION 3. Consider the quadratic cost model and a common grid $k > 0$. The set of coarse grid Nash equilibria corresponding to a grid of size k is (p, p) with $p = kn$ for $n = n^{-}(k), \ldots, n^{+}(k), n^{+}(k)+1, n^{+}(k)+2$.

The CGNE is implemented empirically by estimating k. The econometric model is a mixture of local normal distributions corresponding to *CGNE* points, each with equal probability and each with a distinct variance to be estimated. In the linear cost setting, we take the CGNE set to include the standard NE play as well as k and $2k$. This is done under the behavioral assumption that the value for the marginal cost, $c = 5$, is negligible compared to the price interval [0, 500], and therefore the case (i) in the discussion of proposition 2 applies.[6]

In the quadratic cost setting, we take the CGNE set to be as defined by Equations 5 and 6.

A model with m *CGNE* points thus has $m+1$ parameters to be estimated: the grid size k and the standard deviation for each equilibrium play. The probability that a player of type n who plays a given *NE* type submits price equal to i, conditional on price i being greater than or equal to 0 (due to censoring of the price space), is equal to the standardized normal probability density function with a correction for censoring:

$$\begin{aligned} &\Pr(price = i|type = n) = \frac{1}{\sigma_n}\phi\left(\frac{NE_n - i}{\sigma_n}\right)/(1 - \Phi\left(-\frac{NE_n}{\sigma_n}\right), \ i > 0, \\ &\Pr(price = 0|type = n) = \Phi\left(-\frac{NE_n}{\sigma_n}\right), \ i = 0. \end{aligned} \tag{7}$$

We assume a mixture of equal probability of belonging to each type. Technically, we could let the segment size be a free parameter to be estimated. Mixture models estimation permits this freedom. However, this would hurt both parsimony and robustness, and the theory does not favor one equilibrium over another, so to keep the model simple, we keep the mixture probabilities fixed. However, we do test for the number of equilibria by running the model with different numbers of equilibria (one, two, and three in the linear cost setting; one to five in the quadratic cost case) and comparing fit.

4. Experimental Results

In this section we first present some descriptive statistics, which will help us assess the presence and relevance of the Bertrand paradox in our experiment, both in the standard linear version and in the quadratic game. After that, we present and discuss the estimates of our coarse grid model together with those of the Quantal response and the step-level thinking model. The evolution of average prices shown below across rounds strongly suggests that the evolution of prices gets remarkably flat after a few rounds. The first and last blocks of five

[6] Recall our discussion above.

Table 2. Descriptive Statistics: Posted Prices, FM = Fixed Match; RM = Random Match

	Posted Price (SD)			Median		
Game	All	R01–R05	R16–R20	All	R01–R05	R16–R20
Linear FM	84.94 (90.04)	109.89 (91.03)	68.17 (88.45)	55	90	43.5
Linear RM	82.81 (75.07)	126.11 (92.76)	48.90 (38.02)	62	110	44
Quadratic FM	165.72 (97.47)	176.68 (85.76)	169.46 (119.73)	130	150	115

periods seem natural candidates to test the ability of the different models to fit with different types of data, in the short and long term.[7]

Summary Statistics

Tables 2 and 3 show posted and market prices for the first five periods, the last five periods, and averaged over all rounds. It also displays median prices so as to assess the information conveyed by the average prices.

The first feature is the different success of the Nash concept for organizing data across games: In the linear Bertrand setting with fixed matching, prices are far from the Nash equilibrium (5) in the earlier rounds (the average market price in the first five rounds is 79.26) and stay so along the experiment (the overall average market price is 66.61) despite the fact that prices decline over time (the average market price in the last five rounds is 57.61). For the overall rounds, the median market price is 50, which indicates that prices are kept away from the equilibrium. This pattern does not change when a random matching protocol is used: Although slightly lower average market prices are observed (for example, the overall market price is 58.86), they are well above the Nash equilibrium, as the overall median market price is also 50.

Table 3. Descriptive Statistics: Market Prices, FM = Fixed Match; RM = Random Match

	Market Price (SD)			Median		
Game	All	R01–R05	R16–R20	All	R01–R05	R16–R20
Linear FM	66.61 (70.82)	79.26 (63.65)	57.61 (76.08)	50	65	35
Linear RM	58.86 (44.30)	85.31 (60.99)	39.00 (26.25)	50	80	37
Quadratic FM	148.82 (87.55)	149.91 (68.43)	158.07 (113.57)	119	130	109

[7] As in all experiments games are played 20 times, an alternative approach to our experimental data would be to analyze how players learn to play the game. Our results strongly suggest that learning effects are secondary after the initial rounds. We believe our approach is useful to compare alternative theories in two different environments: when players start playing the game (first five rounds) and at the end of the experiment (in the last five periods). As we show later in this section, the comparison of the different models yields very similar results in both scenarios.

In the quadratic game, market and posted prices are much closer to the Nash prediction [50, 150]. The overall average posted price is 148.82, just below the efficient equilibrium of 150. It is actually quite remarkable that in the first five rounds, the average market price (149.91) matches the efficient equilibrium (150). Note also that the median prices, both in terms of posted and market prices, lie inside the equilibrium prediction interval for all time horizons.

The evolution of posted and markets prices across periods by treatment are depicted in Figure 1. We use vertical box plots with the variable period in the horizontal axis. The main box captures the vast majority of data, as it contains prices in the 25th–75th percentiles with the median price represented by a horizontal line inside the box. Whiskers include adjacent values, and nonadjacent values are represented by dots.

Figure 1 makes apparent that prices in all treatments settle down after a few periods of downward adjustment. For the quadratic case, prices settle down around the efficient NE from the early stages of the game, but for the linear Bertrand games, prices are always well above the NE. The random matching protocol case reduces price dispersion (measured by the height of the box) and eliminates collusive market prices already from early stages of the game. However, no other major differences in the evolution of prices are observed with respect to the fixed matching.

Result 1. Over time, prices move closer to NE in both games. They stay away from equilibrium in the linear Bertrand setting but get remarkably close to the efficient equilibrium in the quadratic game.

Result 2. The main effect of the random matching is that it eliminates highly collusive market prices, but it does not affect average behavior.

We next focus on price distributions in the short and long run. Figures 2 and 3 show the histograms of posted and market prices in the first five and last five rounds, respectively.[8]

The first observation is that prices are diffused in the interval, although there are clear peaks in the distribution, and these spikes correspond to what can be broadly named ‘‘prominent’’ or ‘‘round’’ numbers. Consider, first, the linear Bertrand setting. We already know that the Nash concept predicts poorly. In the first five periods, the Nash price 5 does not stand out; the local modes of the price distribution correspond to prices 100, 150, and 200 (these findings get amplified in the random matching treatment and in market price distribution). It is not until the last five rounds that the Nash prediction stands out as the modal price for the fixed matching (but not for the random matching treatment, where 50 is the most chosen price) although most of the prices fall to the right, with clear spikes at 50 and 100.

For the quadratic game, we also observe peaks at prominent numbers. But contrary to the linear Bertrand setting, the modal price corresponds to a prominent number inside the Nash equilibrium interval. In the first five rounds, the spikes take place at prices 150, 100, 200, and 250. In the last five rounds, price 100 stands out as the modal price, although 150 and 200 are clear peaks. The large peak in the quadratic cost case at a price of 100 may indicate that 100 has a special prominence in this setting, perhaps because 100 is the marginal cost when the firms tie.

Result 3. Most of the data fall to the right of the Nash prediction in both games. The histograms reveal peaks at prominent numbers

---

[8] Note the different vertical scale for the quadratic case, where the dispersion of data is lower.

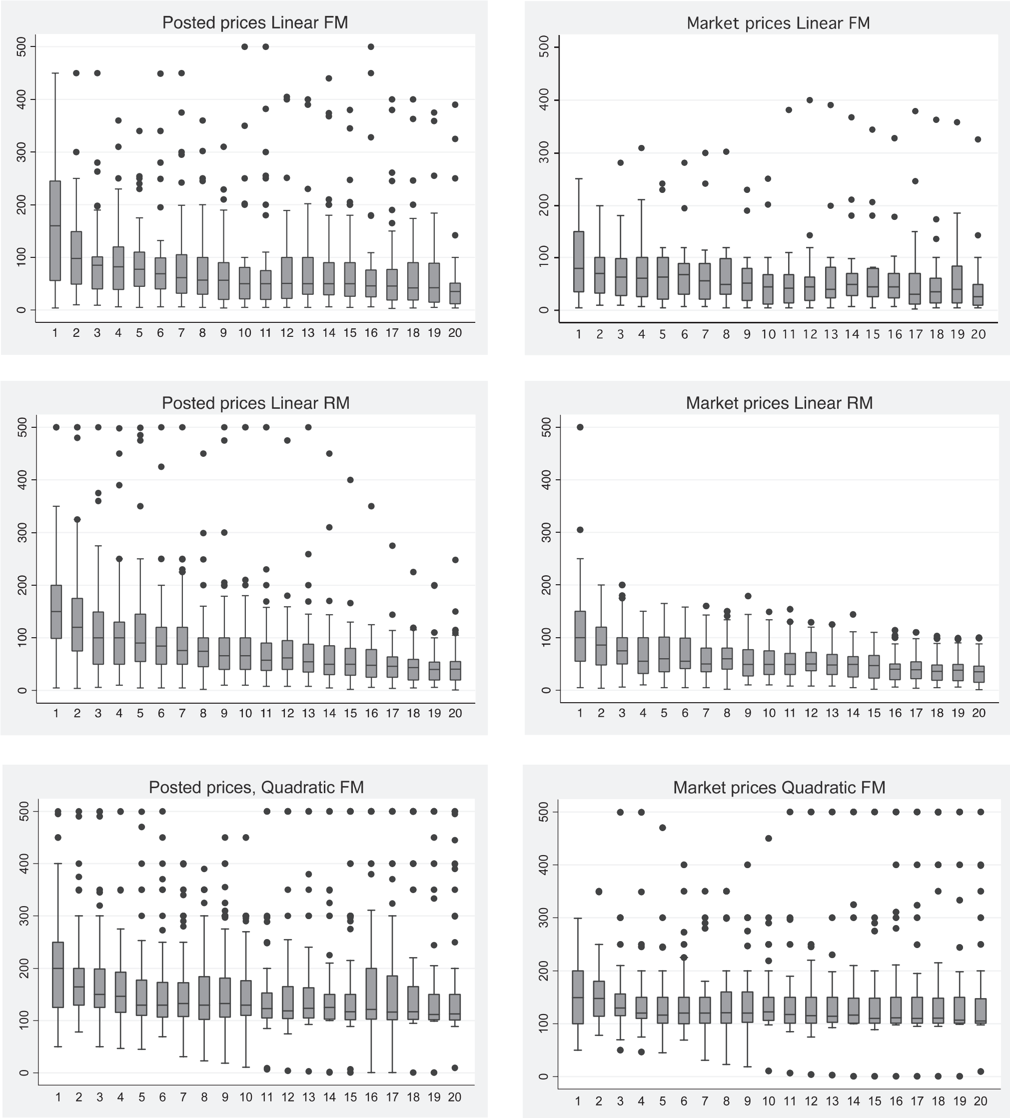

Figure 1. Posted and Market Prices in the Three Treatments

Econometric Estimations

We now present the econometric estimations of the three models we presented: QRE, level-1, and CGNE. The goal of this empirical exercise is simple. We just want to identify whether a behavioral model consistently fits better across all cost specifications and matching conditions. Figures 4 and 5 show the cumulative distributions for the three models we tested: QRE, level-1, and CGNE.

Because of the clearly different grid sizes in the first and last blocks of five periods, we estimated behavior separately for each of these blocks of five periods. We report on the first block of five periods first.

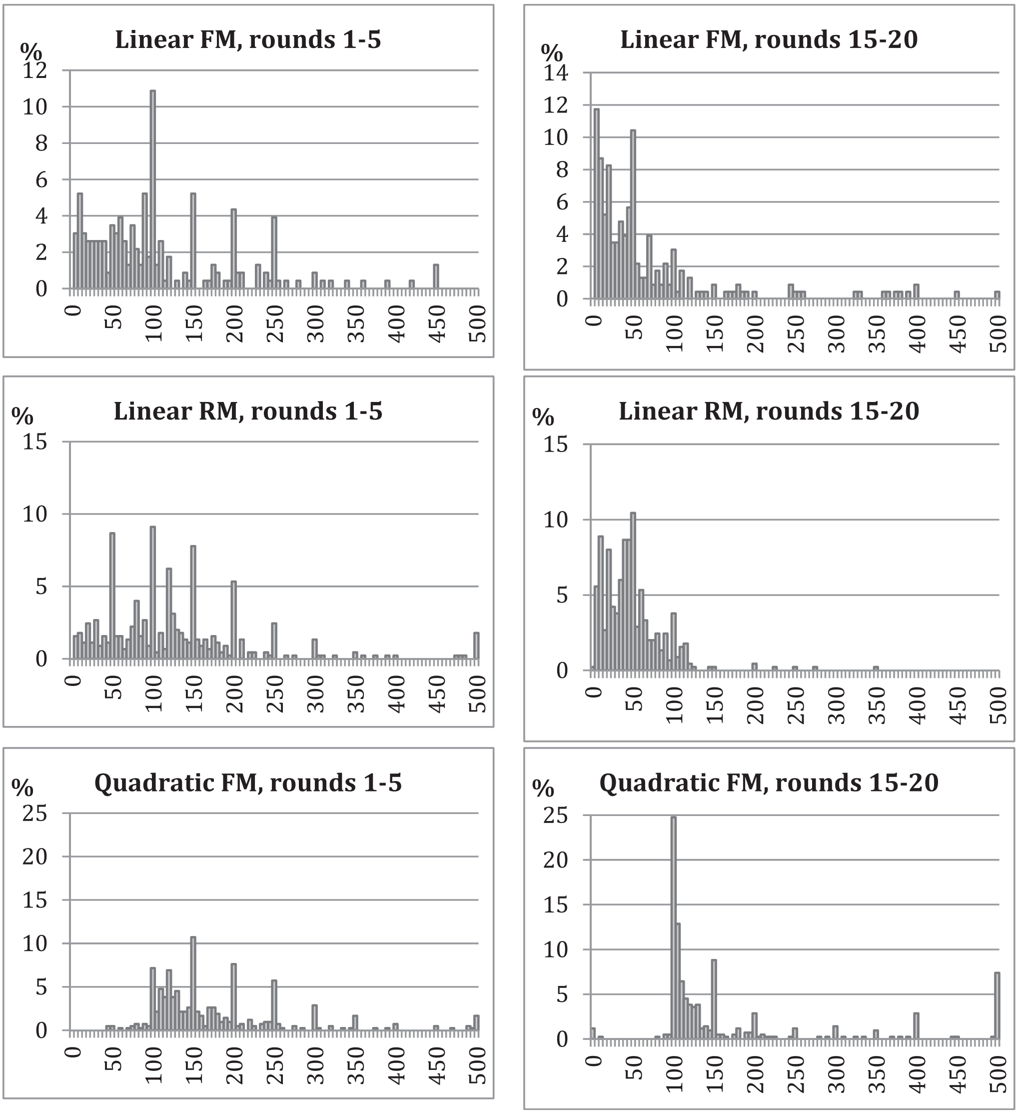

Figure 2. Histograms of Posted Prices

Linear Bertrand Competition

The log likelihoods for the random match Linear Bertrand case, first five periods, are −2619, −2792, and −2577 for QRE, level-1, and CGNE, respectively. The CGNE model produces a better fit than alternative models, with an estimated grid of 103 (standard error of 13), and the equilibria are 5, 103, and 206. The log likelihoods for the fixed match, first five periods, are −1365, −1368, and −1298 for QRE, level-1 and CGNE, respectively. The estimated grid size k for the first five periods is 90 (standard error of 25), and the estimated number of equilibria is three, corresponding to c, k, and $2k$. The model estimation indicates three subpopulations with means at 5, 90, and 180.

For the last five periods, the log likelihoods for the random match Linear Bertrand case are −2300, −2692, and −2170 for QRE, level-1, and CGNE, respectively. The estimated grid

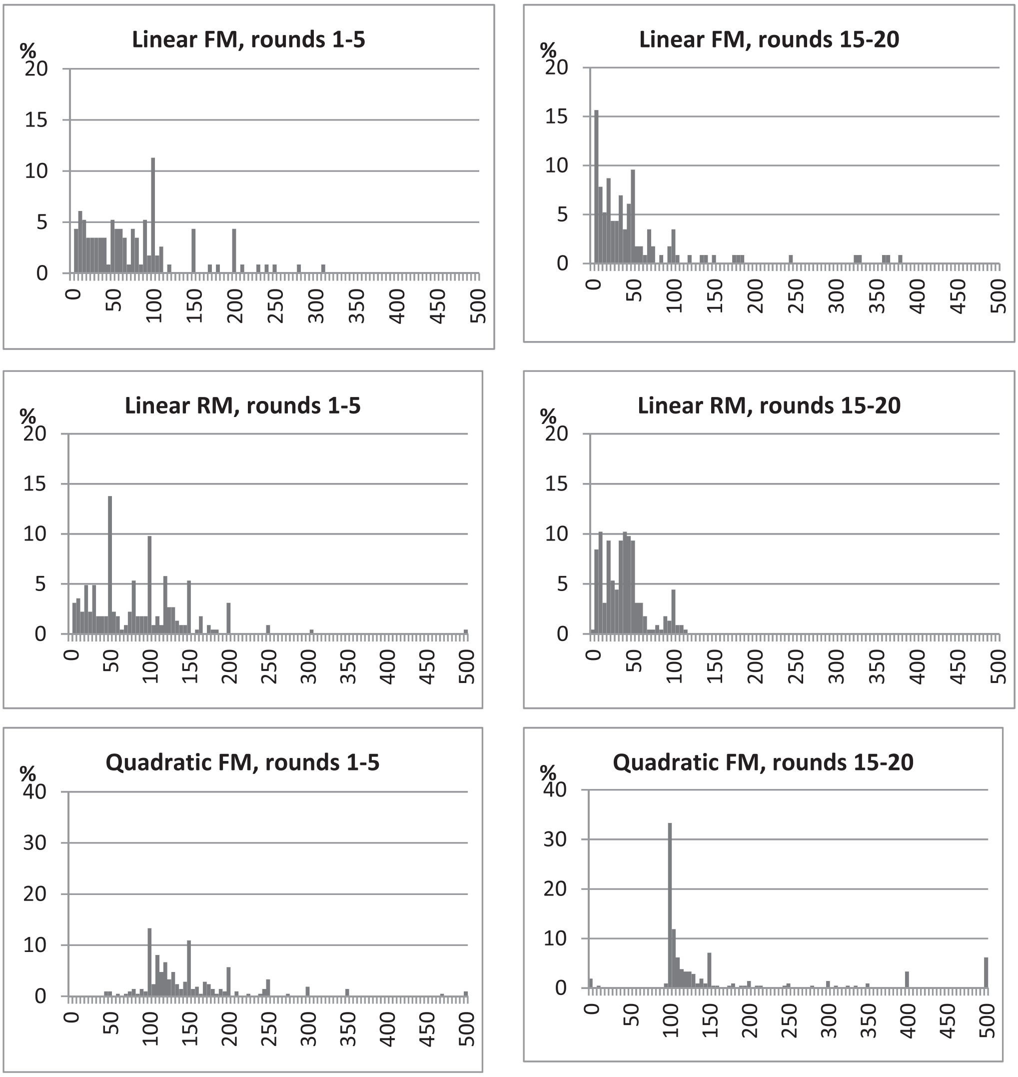

Figure 3. Histograms of Market Prices

size for the last five periods is 34 (standard error of 4), and the equilibria are 5, 34, and 68. For the fixed match case, the log likelihoods for the last five periods are −1404, −1407, and −1187 for QRE, level-1, and CGNE, respectively. Hence, again the CGNE model outperforms alternative models, although in this case, the estimated grid size is 0.[9] This means that it collapses to the standard NE model with noise for the fixed matching condition.

The comparison of the estimated grids for the first five and the last five periods blocks and across matching protocols is interesting. In the first five periods block, the estimated grid is around $k = 100$, in both matching protocols suggesting that subjects chop the price interval in five equidistant subintervals. However, our estimations reveal finer grids for both matching

[9] This happens despite the existence of a clear local mode at 50. The estimated value is 0 because the global mode is around the standard NE and there is not much of a mode at 100.

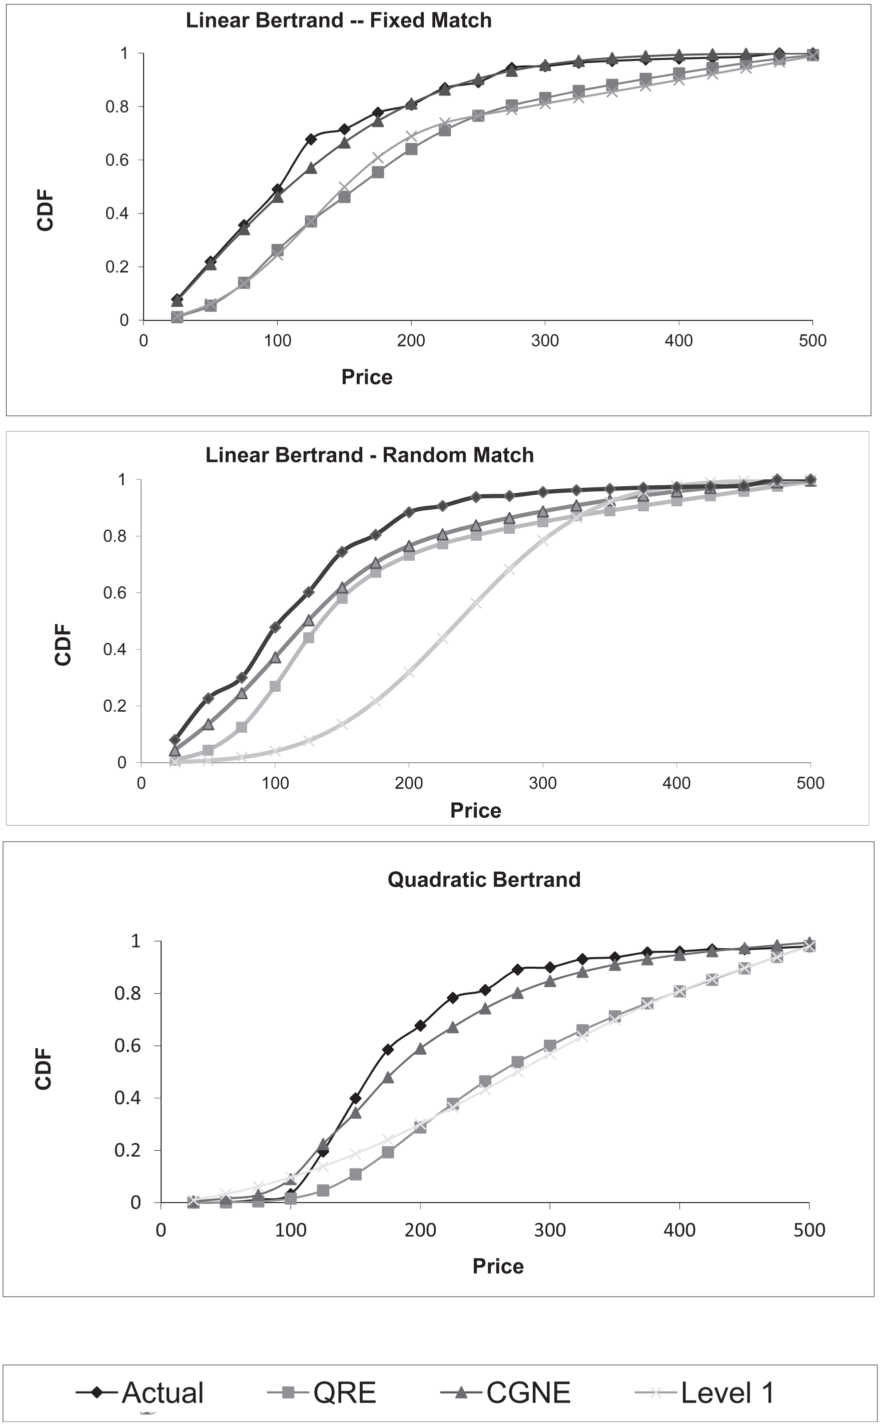

Figure 4. Cumulative Distribution Functions (First Five Rounds)

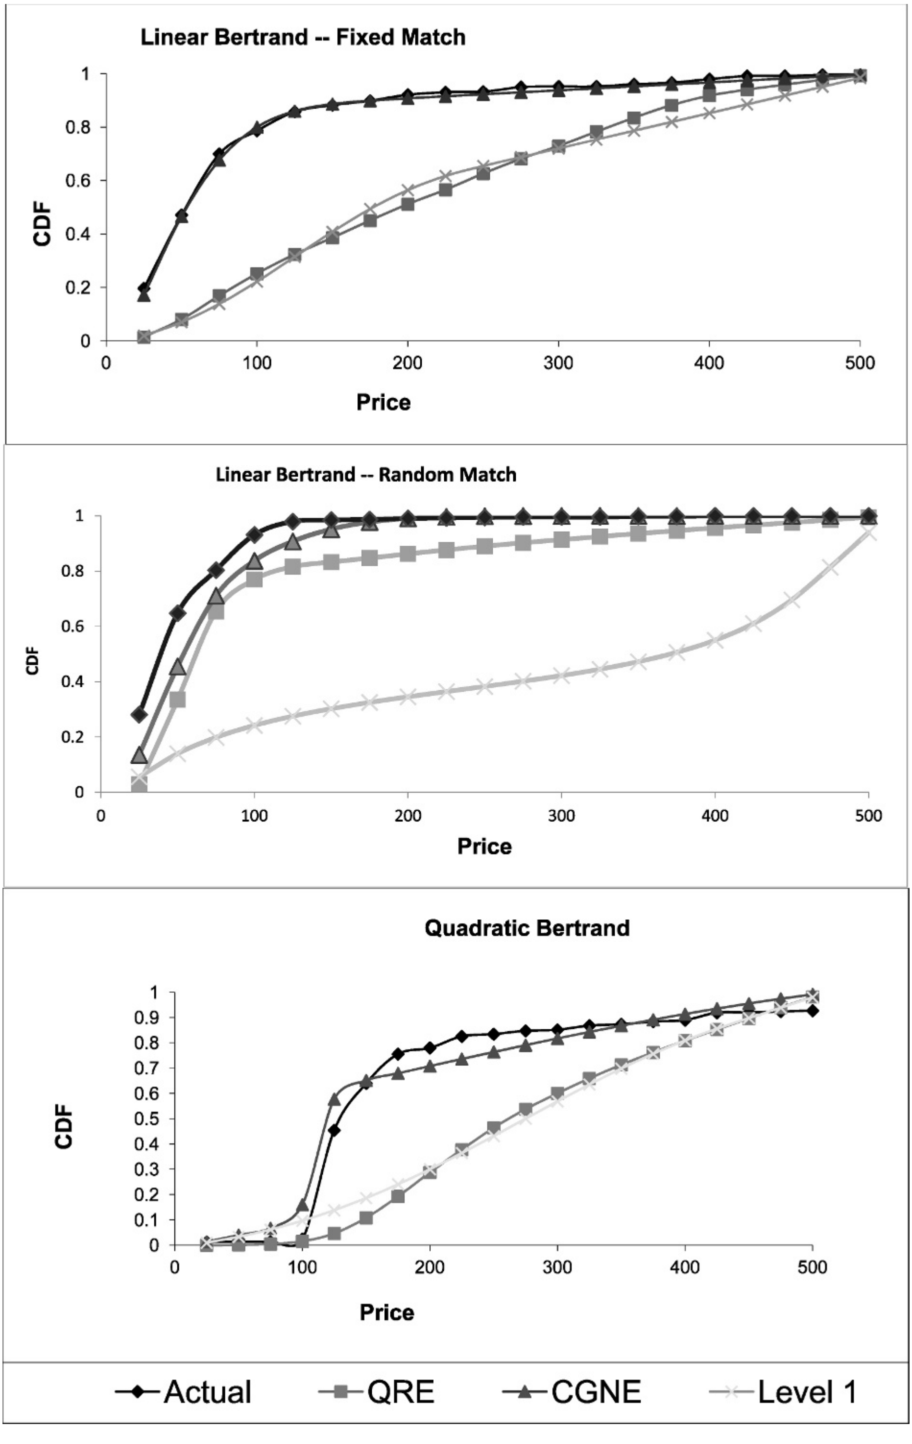

Figure 5. Cumulative Distribution Functions (Last Five Rounds)

protocols in the final part of the experiment. Although our model does not explicitly address the dynamic adjustment of grids, an intuitive explanation would be that competition might extend to grid selection when price competition is exhausted. Naturally, the "exhaustion of price competition" is more likely to happen when competition takes place repeatedly with the same opponent. This is exactly what the econometric estimations reveals.

RESULT 4. In the linear Bertrand game, the coarse grid Nash equilibrium outperforms the competing behavioral explanations. The estimated number of equilibria is three, with means 5, 90, and 180 with fixed matching, and 5, 103, and 206 with random matching. In the last five periods, the equilibrium prediction converges to standard NE in the fixed match setting.

Quadratic Bertrand Competition

In the first five periods of the quadratic cost case, the log likelihoods are −2525, −2589, and −2356 for QRE, level-1, and CGNE, respectively.

The grid parameter k is estimated at 50 and there are four subpopulations identified, with means at 100, 150, 200, and 250. There is no mode at 50. This is not necessarily a problem when we do not force the estimation to begin at the lowest possible mode for a given k. When we do force that restriction, the grid parameter k is estimated to be 75, which captures the mode at 150, and fits the mode at 100 through a higher estimated variance. That restricted model still outperforms the QRE and level-1 models, with a log likelihood of −2388.[10]

In the last five periods of the quadratic cost case, the log likelihoods are −2610 for QRE (with QRE parameter estimated at 0, implying completely and uniformly random choice), −2608 for level-1 (not significantly different from random choice), and −2266 for CGNE with two identified modes at 100 and 150 (−2488 when the model is restricted to begin at the lowest possible mode, still outperforming QRE and level-1). In other words, for the last five periods, the noisy fit of QRE and level-1 does not help in light of the sharp concentration of choice around 100. For the CGNE model, the grid parameter k is estimated at 50, and there are only two subpopulations identified, with means at 100 and 150. The parameter k is not statistically different from the first five periods, but the number of modes is statistically different between the first five periods and the last five periods.[11]

RESULT 5. In the quadratic pricing game, the coarse grid Nash equilibrium outperforms the competing behavioral explanations. For the first five periods, coarse grid Nash equilibria with means 100, 150, 200, and 250 are found in the population. In the last five, only two equilibria are identified at 100 and 150.

Note that although the coarse grid Nash model does not fit the local mode at 200, it does provide a natural interpretation to two observed local modes within the Nash range. This bimodal distribution cannot be explained with a straightforward Nash model.

Level-1 and QRE are nearly indistinguishable in the Linear Bertrand case with fixed matching, whereas QRE outperforms in terms of likelihood Level-1 in the linear case with

[10] There is no mode at 50 because among all equilibrium prices, those lower than 100 are risky because any miscoordination entails negative profits for the low price firm due to the strict convexity of the cost function. The first equilibrium price free from this miscoordination threat is 100, and this explains the local mode at 100 and the absence of posted prices below it (recall Figure 2).

[11] This result is compatible with the previous dynamic picture in that there is no collapse of individual grids, recall that in the quadratic case we had fixed matching, but price competition takes the form of choosing lower CGNE.

random matching and the Quadratic Bertrand case. The coarse grid Nash model greatly outperforms in terms of likelihood the other two in both scenarios, although less clearly with random matching. We believe that the superiority of the CGNE model over the other two is not driven by its greater flexibility—measured by the number of parameters to be estimated—but by its natural ability to generate the multimodal price distributions (as Figures 2 and 3 suggest). 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.

5. Conclusions

Bertrand price competition is a useful abstraction of markets in many settings and one that naturally lends itself to laboratory experiments. There are different prediction approaches in Bertrand settings, but the Nash equilibrium approach has been somewhat discredited in the literature as a predictive tool on the grounds that it lacks the ability to characterize human behavioral patterns. In this work we show that a minor modification to the Nash equilibrium approach can restore its standing as a predictive tool. Specifically, we show that for a 20-period run in the standard two-player Bertrand setting, the coarse grid Nash equilibrium prediction has a high likelihood relative to the other approaches in three different treatments: linear costs with fixed and random matching and quadratic costs with fixed matching.

In the three treatments, we observe a price pattern that displays local modes at convenient round prices from a human perspective: 50, 100, 150, and 200. This pattern is consistent with jump bidding. Our experiments are not designed to identify explanations for jump bidding. We simply observe jump bidding as an empirical pattern and examine the predictions of the Bertrand model when jump bidding follows fixed increments that are larger than the minimal increment. We prove that the equilibrium predictions change, and we aim to see whether the new equilibrium predictions correspond better to the data than models of bounded rationality previously proposed in the literature.

In the linear cost case, and for both the fixed and the random matching, market prices early or late in the game are far from marginal cost. Moreover, they appear to have modes at 50 and 100. The coarse grid Nash model with $k = 50$ would predict such modes. However, we also observe modes at 150 and 200 early in the game. Thus, it appears that different pairs of subjects coordinate on different values of k. Our coarse grid Nash model estimation yields an average k of 90, and this appears to correspond to the observed data. The fit in terms of maximized likelihood and empirical distribution function is far better than that of the alternative models we examined.

For a quadratic-cost case, the pure Nash prediction is within 1 of both the mean market price in the last round and in all 20 rounds. This could be considered good predictive power. Nevertheless, the coarse grid Nash model still adds predictive power in the quadratic case by capturing the bimodal distribution.

Appendix

Proof of Proposition 1. (For the quadratic game.)

1. There are no asymmetric NE. By contradiction, let (p_1^*, p_2^*) be any asymmetric Nash equilibrium and let $p_1^* < p_2^*$. Then

$$\frac{\partial \pi_1}{\partial p_1}(p_1^*, p_2^*) = Q > 0. \tag{A.1}$$

Thus, firm 1 can increase profits by increasing its price. Contradiction.

2. The lower limit of the set of symmetric NE yields zero profit

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

For given price p larger than p^- the profit is

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

The deviation to consider is to set a price $p - \varepsilon$ with $\varepsilon > 0$. The associated profit is

$$\pi_i(p-\varepsilon, p) = (p-\varepsilon)Q - cQ^2, \tag{A.4}$$

and the difference is

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

Hence, for any price lower than $\frac{3}{2}cQ$ there is no profitable deviation.

Proof of Proposition 2. In the linear case, it is trivial that we need to focus on symmetric price configurations. Consider a grid $k > 0$ and n integer. For a price $p = nk$ being an equilibrium one, it is needed that (i) it yields non-negative profits and that (ii) there is no profitable undercutting.

From the first requirement we obtain

$$\pi(nk, nk) \geq 0 \rightarrow (nk - c)\frac{Q}{2} \geq 0 \rightarrow n \geq \frac{c}{k},$$

and from the second one we obtain

$$\pi((n-1)k, nk) \leq \pi(nk, nk) \rightarrow (nk - k - c)Q \leq (nk - c)\frac{Q}{2} \rightarrow n \leq 2 + \frac{c}{k}.$$

Hence, we find $\frac{c}{k} \leq n \leq 2 + \frac{c}{k}$ as required.

Proof of Proposition 3. Assume that subjects use prices which are multiples of k. We proceed in three steps:

Step 1. It is trivial to show that no firm has negative profits in equilibrium.

Step 2. We now show that there are no asymmetric equilibria. Let (p_1^*, p_2^*) be any asymmetric Nash equilibrium. Prices can be written as $p_1^* = n_1 k$ and $p_2^* = n_2 k$. Let $n_1 < n_2$. Then $\pi_1(p_1^*, p_2^*) \geq 0$ and $\pi_2(p_1^*, p_2^*) = 0$. For these prices to be in equilibrium it must be the case that $n_2 - 1 \leq n_1$ because otherwise $n_1 + 1$ would be a profitable deviation for firm 1. Hence, the case we need to analyze is $n_2 = 1 + n_1$.

There are two possibilities:

(a) $\pi_1(p_1^*, p_2^*) = 0$. Firm 1 will deviate by matching firm 2's price.

(b) $\pi_1(p_1^*, p_2^*) > 0$. Firm 2 will deviate by matching firm 1's price.

Step 3. It is easy to see that all multiples of k in the Nash interval $\left[\frac{1}{2}cQ, \frac{3}{2}cQ\right]$ are Nash equilibrium of the game. From Equation A.5, for a rival's price of $p = nk > \frac{3}{2}cQ$, the maximum profitable deviation is $\frac{1}{2}\left(nk - \frac{3}{2}cQ\right)$. We compute the rival's n at $p = nk$ for which the maximum possible deviation is exactly equal to k. The solution is $n = 2 + \frac{3}{2}\frac{cQ}{k}$. Note that $int\left(\frac{3}{2}\frac{cQ}{k}\right)$ is n^+. Thus, $n^+ + 1$ and $n^+ + 2$ are also Nash equilibria.

Instructions for Linear Bertrand Setting, Random Match

1. Welcome. Your show up fee is €5. This experiment has 20 independent rounds. At the beginning of the experiment, you will be assigned to a cohort of 10 participants. In each round, you will be randomly matched with a participant

from your own cohort to form a group (which will be called *market*) of two participants. Markets are independent and you will never be informed of the identity of the participants in your cohort or your market.
2. In the experiment, each participant is a firm making pricing decisions. Prices need to be between 0 and 500 ECUs (the experimental currency unit used in this experiment).
3. The *profits* of your firm are calculated as the *revenues* minus the *costs*. The *revenue* is the product of your demand (the number of units that you sell) multiplied by your selling price (the price at which you sell). Your demand in each round depends exclusively on your decisions (*your price*) and the decisions of the other firm in your market (*the other price*).
4. The market *Demand* in each round is *fixed*, and given that the two firms offer exactly the same product, two possible scenarios can happen:
   a) Either the two prices are equal, in which case the firms equally share the demand.
   b) Or the prices are different, in which case the firm offering the lower price gets the whole market demand in that round (the other firm does not sell anything).
5. The *cost* function is linear, so that the cost of producing the whole market is two times the cost of producing half market. Note that no production is costless and that regardless of your market share, you are required to attend the whole demand.
6. In a given round you can obtain *profits* or *losses*, which will be computed following each round, but if at the end of the experiment your accumulated profits are negative, these losses will never become effective. At the beginning of each round you will know the past values of prices and benefits obtained by the two companies in previous rounds and your accumulated benefits.
7. At the end of the experiment, your accumulated profits in ECUs will be converted to Euros with an exchange rate 4000 ECU = €1.

References

Abbink, Klaus, and Jordi Brandts. 2008. Pricing in Bertrand competition with increasing marginal costs. *Games and Economic Behavior* 63:1–31.

Apesteguia, Jose, Steffen Huck, and Jörg Oechssler. 2007. Imitation—Theory and experimental evidence. *Journal of Economic Theory* 136:217–235.

Apesteguia, Jose, Steffen Huck, Jörg Oechssler, and Simon Weidenholzer. 2010. Imitation and the evolution of Walrasian behavior: Theoretically fragile but behaviorally robust. *Journal of Economic Theory* 145:1603–1617.

Argenton, Cedric, and Wieland Müller. 2012. Collusion in experimental Bertrand duopolies with convex costs: The role of cost asymmetry. *International Journal of Industrial Organization* 30:508–517.

Ascioglu, Asli, Carole Comerton-Forde, and Thomas H. Mclnish. 2007. Price clustering on the Tokyo stock exchange. *Financial Review* 42:289–301.

Baye, Michael R., and John Morgan. 2004. Price dispersion in the lab and the internet: Theory and evidence. *RAND Journal of Economics* 35:449–466.

Chamberlin, Edward H. 1948. An experimental imperfect market. *Journal of Political Economy*. 56:95–108.

Dastidar, Krishnendu G. 1995. On the existence of pure strategy Bertrand equilibrium. *Economic Theory* 5:19–32.

Dufwenberg, Martin, and Uri Gneezy. 2000. Price competition and market concentration: An experimental study. *International Journal of Industrial Organisation* 18:7–22.

Edgeworth, Francis Y. 1925. The pure theory of monopoly. In *Papers Relating to Political Economy*. London: MacMillan. 1, 111–42.

Fatas, Enrique, and Antonio J. Morales. 2013. Step thinking and costly coordination. *Economics Letters* 120:181–183.

Fouraker, Lawrence, and Sidney Siegel. 1963. *Bargaining behavior*. New York: McGraw-Hill.

Friedman, L. 1967. Psychological pricing in the food industry. In *Prices: Issues in theory, practice, and public policy*, edited by Almarian Phillips and Oliver E. Williamson. Philadelphia: University of Pennsylvania Press, pp. 187–201.

Gigerenzer, Gerd, Peter M. Todd, the ABC Research Group. 1999. *Simple heuristics that make us smart*. New York: Oxford University Press.

Gneezy, Uri. 2005. Step-level reasoning and bidding in auctions. *Management Science* 51:1633–1642.

Harris, Lawrence. 1991. Stock price and discreteness. *Review of Financial Studies* 4:389–415.

Isaac, R. Mark, Timothy C. Salmon, and Artie Zillante. 2007. A theory of jump bidding in ascending auctions. *Journal of Economic Behavior & Organization* 62:144–164.

Janssen, Maarten, and Eric B. Rasmusen. 2002. Bertrand competition under uncertainty. *Journal of Industrial Economics* 50:11–21.

Jehiel, Philippe. 2005. Analogy-based expectation equilibrium. *Journal of Economic Theory* 123:81–104.

Kreps, David M., and Jose A. Scheinkman. 1983. Quantity precommitment and Bertrand competition yields Cournot outcomes. *Bell Journal of Economics* 14:326–337.

Kwasnica, Anthony M., and Elena Katok. 2007. The effect of timing on bid increments in ascending auctions. *Production and Operations Management* 16:483–494.

Mitchell, Jason. 2001. Clustering and psychological barriers: The importance of numbers. *Journal of Futures Markets* 21:395–428.

Neugebauer, Tibor, and Reinhard 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.

Niederhoffer, Victor. 1965. Clustering of stock prices. *Operations Research* 13:258–265.

Niederhoffer, Victor. 1966. A new look at clustering of market prices. *Journal of Business* 39:309–313.

Osborne, M. F. M. 1965. The dynamics of stock trading. *Econometrica* 33:88–111.

Payne, John W., James R. Battman, and Eric J. Johnson. 1988. Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory and Cognition* 14:534–552.

Rubinstein, Ariel. 1998. *Modeling bounded rationality*. Cambridge, MA: MIT Press.

Schindler, Robert M., and Patrick N. Kirby. 1997. Patterns of rightmost digits used in advertised prices: Implications for nine ending effects. *Journal of Consumer Research* 24:192–201.

Singh, Nirvikar, and Xavier Vives. 1984. Price and quantity competition in a differentiated duopoly. *RAND Journal of Economics* 15:546–54.

Sonnemans, Joep. 2006. Price clustering and natural resistance points in the Dutch stock market: A natural experiment. *European Economic Review* 50:1937–1950.

Stahl, Dale O., and Paul Wilson. 1994. Experimental evidence of players' models of other players. *Journal of Economic Behavior and Organization* 25:309–327.

VanRullen, Ruffin, and Christof Koch. 2003. Is perception discrete or continuous? *TRENDS in Cognitive Sciences* 7:207–213.

Whynes, David K., Zoë Phillips, and Emma Frew. 2005. Think of a number … any number? *Health Economics* 14:1191–1195.