Bertrand Competition Under Uncertainty

June 24, 1998

Maarten Janssen and Eric Rasmusen

Abstract

Consider a Bertrand model in which each firm may be inactive with a known probability, so the number of active firms is uncertain. This simple model has a mixed-strategy equilibrium in which industry profits are positive and decline with the number of firms, the same features which make the Cournot model attractive. Unlike in a Cournot model with similar incomplete information, Bertrand profits always increase in the probability other firms are inactive. Profits decline more sharply than in the Cournot model, and the pattern is similar to that found by Bresnahan & Reiss (1991).

Janssen: Department of Economics, H09-20, Erasmus University, 3000 DR Rotterdam, The Netherlands. Fax: 31-10-2121724. Email: Janssen@micro.few.eur.nl.

Rasmusen: Indiana University, Kelley School of Business, BU 456, 1309 E 10th Street, Bloomington, Indiana, 47405-1701. Office: (812) 855-9219. Fax: 812-855-3354. Email: Erasmuse@Indiana.edu. Web: Php.indiana.edu/~erasmuse. This paper is available on the web in Acrobat format at Php.indiana.edu/~erasmuse/@Articles/Unpublished/bertrand.pdf.

We thank for their comments David Schmidt and seminars at the Indiana University Dept. of Economics and CIRANO in Montreal.

Footnotes beginning with xxx are notes to ourselves for redrafting.

1. Introduction

Consider a carpenter who is asked by a house owner to submit a tender for renovating part of the house. He considers it very likely that if the homeowner has asked for tenders from other carpenters, he gives the job to the carpenter with the lowest tender. However, he also thinks there is a chance that the homeowner has not asked anybody else for a tender and will give the job to him if only he does not ask more than the owner is willing to pay. What price will the carpenter charge?

We try to answer that question. With some probability the carpenter is a monopolist who can charge the monopoly price but with some probability he faces Bertrand competition. We will show that there exists a unique Nash equilibrium, in mixed strategies. As the expected industry profits are positive for any number of firms N, the model does not suffer from weaknesses of the standard Bertrand model. Moreover, the model does reasonably well in explaining the empirical results of Bresnahan and Reiss (1991) on *how* industry profits decline with increasing N.

The model allows for many interpretations, in two categories. First, uncertainty about the existence of competitors may arise from uncertainty about consumer search behavior, as with the carpenter above. It may be unclear whether consumers regard rival commodities as perfect substitutes, consumer search costs may be uncertain from the firm's perspective, and consumers may vary in their sophistication.

Second, uncertainty about the existence of competitors may arise from uncertainty about other sellers' behavior. It may be unclear whether rivals have hit their capacity constraints (in which case they cannot compete for additional consumers), whether rivals have entered yet, whether rivals have grossly overpriced by mistake or ignorance, or whether rivals have temporarily high costs. It may be unclear whether other competitors have also discovered a new market, or in black markets it may be difficult to know the number of firms operating in that market (cf., Janssen and Van Reeven, 1998). Any of these situations can be modelled as uncertainty over the number of active rivals.

The paper is related to several different literatures. The paper most closely related is Burdett and Judd (1983). They show that equilibria with price dispersion exist in competitive markets, i.e., different firms may charge different prices, if consumer search is noisy (some consumers observe only one price, others two, and so on). Our model differs in a number of respects. First, while one possible interpretation of our model is that consumers differ with respect to the number of prices they observe, our model allows for many other interpretation (see above) which do not fit the model analyzed by Burdett and Judd (1983). Second, even if we take differences in the number of observed prices as the interpretation of our model, there are important discrepancies between their models and ours. For example, we analyze a strategic model instead of a competitive one and when analyzing the impact of the number of firms on the market outcome, we assume that demand is given, instead of having a constant measure of consumers per firm.

Spulber (1995) analyzes a model of Bertrand competition when firms' cost functions are private information. He shows that the model has a unique pure strategy equilibrium in which firms set prices above marginal cost and have positive expected profits. In contrast, the firms in our model do not know how many competitors they have, but assume that all competitors (if any) have the same cost structure as they themselves have. Even though the type of uncertainty between Spulber's model and ours is of a vary different nature, the properties of the market equilibrium are to a considerable extent similar: firms set prices above marginal cost and receive positive expected profits.

Finally, our model is also of interest for students of auctions. The similarities between Bertrand price competition and first-price sealed-bid auctions is well-known.[1] Our paper can be regarded as answering the question what is the optimal bid if the number of participants to such a sealed-bid auctions is unknown, as is often the case in procurement bids.

Section 2 of the paper lays out the basic model and solves for the mixed strategy equilibrium. Section 3 compares the outcome in the model with that of a Cournot model, and compares the expected industry profits in our model for different values of N with empirical results obtained by Bresnahan and Reiss (1991). Section 4 shows how the basic model can be generalized to cover some of the alternative interpretations alluded to above. Section 5 discusses a two-period model that endogenizes firms' decisions to compete or not. The mixed strategy equilibrium of this model can be seen as a justification for the assumption that firms do not know how many competitors they have. Section 6 concludes.

2. The Basic Model

Let there be N firms that might produce homogeneous goods in a given market. The firms do not know of each other whether they really produce for the same market or not, i.e., they do not know whether or not they are active competitors. The probability each of them assigns to each of the other firms being an active competitor in the same market is α, where $0 < \alpha < 1$. If the firms are active competitors, the market is described by the Bertrand model of price competition. If, however, there is only one firm active in the market, that firm is a monopolist and can charge the monopoly price. For simplicity, we further assume that there is one consumer, who buys at most one unit of the good, and his maximum willingness to pay for the good is v. In case of tied prices, the consumer picks a firm randomly. Marginal cost is normalized to 0.

First, let us establish that there is no Nash equilibrium with any firm putting

[1] See for example, Baye & Morgan (1997a, b).

positive probability weight on choosing any particular price. Suppose Firm 1 (without loss of generality) charges price p' with positive probability, rather than mixing over a continuous range of prices and putting infinitesimal probability on each. Putting positive probability on a price of zero is not profit maximizing, because if the firm charged the monopoly price of v instead on those occasions, it would have an expected payoff of $(1-\alpha)^{N-1}v$, so let us focus on $p' > 0$.

With some probability, Firms 1 and 2 will be the only firms active in the market. Firm 2 has one of two best responses. The first is to put positive probability on price $p' - \epsilon$ for some small ϵ, since then whenever Firm 1 charges p', Firm 2 will capture the consumer and earn a positive profit. This cannot be an equilibrium, however, since by the same logic Firm 1 can profitably deviate to $p' - 2\epsilon$. Firm 2's second possible best response is to put positive probability on the monopoly price of v. (Firm 2 would do this if p' were close to zero and Firm 1 would heavy weight on it.) But if Firm 2 does this, Firm 1 would deviate by switching from p' to $v - \epsilon$, and the equilibrium breaks down.

Firm 2 does not know that Firms 1 and 2 will be the only active firms, but that will occur with strictly positive probability. Hence, with some positive probability, if Firm 1 is putting positive probability on p', Firm 2 does best by assigning positive probabilities in the ways described in the previous paragraph. Since that paragraph showed that Firm 1 would then deviate from its hypothesized strategy, we have shown that no Nash equilibrium has a firm putting positive probability on any particular price.

Let us therefore construct an equilibrium in mixed strategies with the strategies having a continuous support. Let $F_i(p)$ be the probability that firm i charges a price that is smaller than p. Similarly, let $F_{-i}(p)$ denote the vector of cumulative mixed strategies of all firms except firm i. The expected payoff to firm i of charging a price p_i when all other firms choose a mixed strategy according to $F(p)$ is

$$\pi(p_i, F_i(p)) = \Sigma_{k=0}^{N} \begin{pmatrix} N \\ k \end{pmatrix} (1-\alpha)^k [\alpha(1-F(p_i))]^{N-k} p_i. \tag{1}$$

This expression can be explained in the following way. The probability that exactly $N-k$ out of N firms exist is equal to $\begin{pmatrix} N \\ k \end{pmatrix} (1-\alpha)^k \alpha^{N-k}$. The expected payoff to firm i when exactly $N-k$ firms exist and when it charges a price of p_i is equal to p_i times the probability that each of these $N-k$ firms charge a price that is larger than p_i : $(1-F(p_i))^{N-k}p_i$. Multiplying these two terms and summing up over all k gives the expression above.

Expression (1) is, of course, nothing but an application of the Binomial Theorem, and a standard result says that $\Sigma_{k=0}^{N} \begin{pmatrix} N \\ k \end{pmatrix} a^k b^{N-k} = (a+b)^N$. Applying this to the present case implies that

$$\pi(p_i, F_i(p)) = (1 - \alpha F(p_i))^N p_i. \tag{2}$$

In equilibrium, firm i must be indifferent between all pure strategies that are in the support of the mixed strategy distribution. Hence, it must be that on some interval of prices the derivative of expression (2) with respect to p_i equals zero. Thus, a necessary condition for any equilibrium in continuous mixed strategies is

$$[1 - \alpha F(p_i)]^N - N[1 - \alpha F(p_i)]^{N-1}\alpha f(p_i)p_i = 0, \tag{3}$$

or

$$1 - \alpha F(p_i) - \alpha N f(p_i)p_i = 0, \tag{4}$$

where f is the density function associated with cumulative distribution function F.

It is a matter of straightforward calculations to show that the solution to differential equation (4) is

$$F(p_i) = \frac{1 - (1-\alpha)\sqrt[N]{v/p}}{\alpha}, \tag{5}$$

for $(1-\alpha)^N v \leq p_i \leq v$.

Result (5) implies that there is a unique symmetric equilibrium with continuous support, and we have shown earlier that an equilibrium in pure strategies does not exist The result is summarized in Proposition 1.

Proposition 1. The equilibrium of the Bertrand model with an uncertain number of competitors is in mixed strategies and the distribution function of a player's strategy is given by

$$F(p_i) = \begin{cases} 0 & for \quad p_i \leq (1-\alpha)^N v \\ \frac{1-(1-\alpha)\sqrt[N]{v/p_i}}{\alpha} & for \quad (1-\alpha)^N v \leq p_i \leq v \\ 1 & for \quad p_i \geq v \end{cases} \tag{6}$$

In Figure 1, the cumulative density is depicted for different values of N, using equation (6), with $\alpha = .2$ and $v = 10$ (prices are at intervals of 1, connected). As N increases, each firm chooses relatively low prices with higher probability. As N becomes large, the cumulative density function approaches 1 for all values of p that are strictly positive. Of course, the equilibrium price under perfect competition is also equal to 0. The perfectly competitive outcome can be regarded as the limit case of the present model when the number of firms becomes very large.

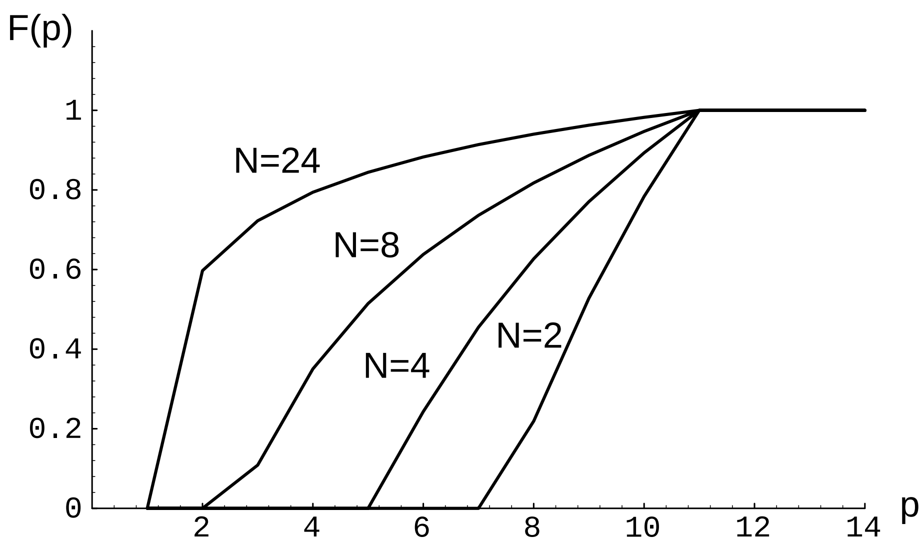

Figure 1: Equilibrium Price Distributions as Industry Concentration Rises

The intuition is straightforward. As the number of potential competitors increases, the probability of at least one other firm actively producing the same product rises. Accordingly, to make profits, a firm must lower its price. In the limit, each firm considers it extremely likely that there is at least one other active competitor. Bertrand competition comes into effect and the resulting equilibrium is such that each firm charges a price that is equal to the perfectly competitive price.

Expected profit for one firm can be found using the pure strategy profit from charging $p = v$, and is, since the firm is active with probability α,

$$\pi_i = \alpha(1-\alpha)^{N-1}v. \tag{7}$$

Note that individual profit is declining in N and its sum[2] is equal to

$$N\alpha(1-\alpha)^{N-1}v, \tag{8}$$

Let π_a denote expected industry profit given that at least one firm is active. The profit in equation (8) can be written as

$$\sum_{i=1}^{N}\pi_i = N\alpha(1-\alpha)^{N-1}v = (1-\alpha)^N(0) + [1-(1-\alpha)^N]\pi_a, \tag{9}$$

[2]Note that although the profits of the different firms are not independent, the expected profits are, so this summation is legitimate.

yielding

$$\pi_a = \frac{N\alpha(1-\alpha)^{N-1}v}{1-(1-\alpha)^N}. \tag{10}$$

To see how industry profit changes with N, note that after some algebra,

$$\frac{d\pi_a}{dN} = \left[\frac{(1-(1-\alpha)^N) + Nlog(1-\alpha)}{(1-(1-\alpha)^N)^2}\right]\left[\alpha(1-\alpha)^{N-1}v\right] \tag{11}$$

a derivative which is well-defined even though only integer values of N have an economic interpretation. The sign of expression (11) is the sign of

$$1-(1-\alpha)^N + Nlog(1-\alpha). \tag{12}$$

For $N = 1$, expression (12) becomes $\alpha + log(1-\alpha)$, which is negative because $\alpha < 1$. For larger N, expression (12) becomes even more negative, because its derivative is $-(1-\alpha)^N log(1-\alpha) + log(1-\alpha) = log(1-\alpha)[1-(1-\alpha)^N] < 0$. Thus,

$$\frac{d\pi_a}{dN} < 0, \tag{13}$$

and profits fall as the number of firms increases.

We can say more. The second derivative can be written, after several steps of algebra,[3] as

$$\frac{d^2\pi_a}{dN^2} = \frac{(1-\alpha)^{N-1}log(1-\alpha)}{[1-(1-\alpha)^N]^2}\left\{1 + \frac{Nlog(1-\alpha)[1+(1-\alpha)^N]}{1-(1-\alpha)^N}\right\}. \tag{14}$$

The first term of this last expression is negative because $log(1-\alpha)$ is negative. The second term has the same sign as

$$1-(1-\alpha)^N + Nlog(1-\alpha)[1+(1-\alpha)^N]. \tag{15}$$

We already found that for all N, expression (12) is negative, i.e.,

$$1-(1-\alpha)^N + Nlog(1-\alpha) < 0. \tag{16}$$

The expression in (16) is negative, and $[1+(1-\alpha)^N > 1$, so $Nlog(1-\alpha)[1+(1-\alpha)^N$ is more negative than $Nlog(1-\alpha)$, so expression (15) must be negative if inequality (16) is true. Hence, since both terms of the last expression in (14) are negative, their product must be positive, and we can conclude that

$$\frac{d^2\pi_a}{dN^2} > 0. \tag{17}$$

[3]xxx This algebra is shown at the end of the working paper as an appendix for referees and other interested parties.

This means that profits are convexly decreasing in the number of firms in the industry, so the shape shown in the numerical examples graphed in Figure 1 would be found for any example.

3. Extensions and Interpretations

The analysis of Section 2 is quite robust to alternative specifications. In particular, we can easily extend the argument to (i) more general demand structures, (ii) different product qualities, and (iii) sequential pricing.[4]

(i) More general demand structures

In Section 2, demand was assumed to be equal to one unit for all prices smaller than v and zero otherwise. In Section 3, demand was assumed to be linear. Here, we will consider a more general demand function, which we denote by $D(p)$. For simplicity we will restrict ourselves to the case $N = 2$. We will impose one condition on this demand function, namely that $pD(p)$ is increasing in p for $p \leq p_m$, where p_m is the monopoly price. Most demand function that are commonly employed satisfy this condition. It is satisfied, for example, if $pD(p)$ is concave in p.

Assumption 4.1. The function $pD(p)$ is increasing and differentiable on $[0, p_m)$.

For general demand functions, the expected profit of firm 1 when firm 2 chooses a price according to the cumulative mixed strategy distribution $F_2(p)$ is given by

$$\pi(p_1, F_2(p)) = (1-\alpha)p_1 D(p_1) + \alpha(1 - F_2(p))p_1 D(p_1). \tag{18}$$

A necessary condition for an equilibrium in mixed strategies with continuous support to exist is that on a certain domain of prices

$$[(1-\alpha) + \alpha(1 - F_2(p_1))][D(p_1) + p_1 D'(p_1)] - f_2(p_1)p_1 D(p_1) = 0. \tag{19}$$

One can show that the solution to this differential equation is given by

[4] One extension we will not attempt here is to risk aversion. The model can, however, illustrate in a simple way the intuition for the result of McAfee and Macmillan (1987) that the seller in an auction can benefit from the risk aversion of the buyers and their lack of precise knowledge of how many bidders are active. In the present model, a seller who charges below B, the lower bound of the support of the mixing distribution, will win the customer with certainty, and still earn a profit. A risk-averse seller would wish to take advantage of this, and would tend to push down the prices charged. A high price is a gamble in the hope that other firms are inactive or are themselves charging high prices, so risk aversion tends to reduce prices.

$$F_2(p) = \begin{cases} 0 & if \quad p \leq \underline{p} \\ \frac{1}{\alpha}\left[1 - \frac{(1-\alpha)p_m D(p_m)}{pD(p)}\right] & if \quad \underline{p} < p \leq p_m \\ 1 & if \quad p > p_m \end{cases} \tag{20}$$

A similar solution holds for Firm 1. It is clear that (20) is similar to (20) and the results of the basic model generalize to more general demand functions. Note that from the solution for $F_i(p)$ it is clear why we have to impose a condition on demand: A necessary and sufficient condition for $F_i(p)$ to be increasing in p is that $pD(p)$ is increasing in p for all values of p smaller than p_m . In the present case it is impossible to provide an explicit solution for the domain of prices over which a firm randomizes. It is clear that the upper bound is given by p_m. This is because even if the other firm does not exist, it is not optimal to set a higher price. The lower bound of the domain, denoted by $\underline{p}$, is defined implicitly by the condition $\underline{p}D(\underline{p}) = (1-\alpha)p_m D(p_m)$ As $pD(p)$ is increasing in p for $p < p_m$, $\underline{p}$ is uniquely defined in this way. Industry profits may be calculated as in the previous section and equal

$$\pi_a = \frac{N\alpha(1-\alpha)^{N-1} p_m D(p_m)}{1-(1-\alpha)^N}. \tag{21}$$

(ii) Different product qualities

So far we have considered the situation that two (or more) firms do not know how many competitors they have. An alternative interpretation of the basic model is that the firms do not know whether consumers consider the products they produce as perfect substitutes or not. To focus ideas and to show that the basic model can be used as a building block for models that analyze different topics in industrial organization, let us suppose two firms can produce different qualities of the same product. They know the quality they themselves produce, but they do not know the quality the competitor produces. In addition, the consumer also does not know which quality the firms produce. For simplicity, let us consider the case in which producers can produce either low (L) or high (H) quality. For any of the two firms, the following is true: consumers and the other firm believe with probability α that the firm produces low quality. As the production capacity is already installed, firms are assumed not to make any choice regarding the quality of their product. Low and high quality incur constant marginal production costs of c_L and c_H, with $c_H > c_L > 0$. Firms only choose prices.

Let us consider the case in which a consumer buys one unit of a product: if he buys the high quality, he does not buy the low quality and vice versa. He derives utility of $v_i - p$ from consuming quality $i, i = L, H$, where p is the price he pays for the product he buys. The consumer maximizes utility. In particular, if the consumer somehow can infer the quality of the products he he will buy the high quality good if and only if $v_H - p_H > v_L - p_L$, where p_H and p_L are the prices charged by the two types of firms. For convenience, assume that $v_L > c_H$. A special case of the present model, is one in which $v_H = v_L$. This case can be interpreted as Bertrand competition with firms not knowing of each other whether they have low or high marginal cost (cf., Spulber, 1995).

There are several types of equilibria in the present model. We mainly consider the case $c_H - c_L > 2(v_H - v_L)$ and show that a continuum of separating equilibria exists, each of which has a structure similar to the one considered in the previous models. As it is not the purpose of the present paper to give a full characterization of the possible equilibrium configurations of the present model we just mention that in case $c_H - c_L < 2(v_H - v_L)$, the equilibrium may be separating, partially separating, or pooling, depending on the parameters.

For convenience, we consider the following two restrictions on the out-of-equilibrium beliefs. First, if a consumer observes an out-of-equilibrium price below c_H, he knows that this price is set by a low-quality firm, because a high-quality firm that maximizes profits would never set such a price. Thus, we impose the belief that $P(\theta = L|p) = 1$ for all $p < c_H$. (Note that this is implied by the Intuitive Criterion of Cho & Kreps (1987) Second, if a consumer observes an out-of-equilibrium price above c_H , it does not know whether it is set by a high or a low-quality firm. It seems reasonable, however, to impose that the consumer has the same beliefs no matter which price above c_H is observed, i.e., $P(\theta = L|p) = \alpha'$ for all $p > c_H$.

Proposition 2. Suppose $c_H - c_L > 2(v_H - v_L)$. For any given level of $\alpha' > 0$, there exists a continuum of equilibria, indexed by θ, where $\theta \in [0, \alpha'(v_H - v_L))$, in which

$$\begin{aligned} p_H &= c_H + \theta \\ F(p_L) &= \tfrac{1}{\alpha}\left[1 - (1-\alpha)\left(\tfrac{(c_H+\theta-c_L)-(v_H-v_L)}{p_L-c_L}\right)\right] \end{aligned} \tag{22}$$

for $(1-\alpha)[c_H + \theta - (v_H - v_L)] \leq p_L \leq c_H + \theta - (v_H - v_L)$.

Proof. First, we calculate the expected profit of the low-quality firm in equilibrium. When this firm sets a price $p_L \leq c_H + \theta - (v_H - v_L)$ it is certain that the consumer will buy from it when the other firm produces high quality. Hence, its expected profit is

$$\pi(p_L|F_L(p)) = (1-\alpha)[(p_L - c_L) + \alpha(1 - F_L(p_L))(p_L - c_L)]. \tag{23}$$

Substituting the expression for $F_L(p_L)$ given in Proposition 2 yields

$$\pi(p_L|F_L(p)) = (1-\alpha)[(c_H + \theta - c_L) - (v_H - v_L)], \tag{24}$$

which is strictly positive.

There are several ways in which the low-quality firm could deviate.

(a) Deviating to a price smaller than $(1-\alpha)[c_H+\theta-(v_H-v_L)]$ makes the consumer buying the product for sure, but it yields strictly smaller profits.

(b) Deviating to a price in between $[c_H+\theta-(v_H-v_L)]$and c_H is not profitable, because the consumer will infer that it is a low quality product and will not buy it.

(c) The firm could deviate by setting p_L in the interval $[c_H, c_H+\theta)$. Given the out-of-equilibrium beliefs we specified, the consumer thinks that there is a probability of a that the good is of a low quality. Thus, the payoff the consumer derives from buying at this price is $\alpha' v_L+(1-\alpha')v_H-p_L$. If the other firm is a high-quality firm, it would set a price of $c_H+\theta$ and the consumer derives a payoff of $v_H-c_H-\theta$ from buying at this price. As θ is smaller than $\alpha'(v_H-v_L)$, the last expression is always larger than the first. Similarly, if the other firm happens to produce low quality, the maximum price it will charge is $c_H+\theta-(v_H-v_L)$ and the consumer's payoff from buying at such prices is larger than or equal to $v_L-[c_H+\theta-v_H-v_L)]$, which reduces to $v_H-c_H-\theta$. Thus, the consumer is always better off buying from the other firm and will not buy at a price in the interval $[c_H, c_H+\theta)$. Hence, it is not beneficial to deviate to such a price.

(d) The low quality firm could deviate by setting its price equal to $c_H+\theta$. The consumer will think that the price is set by a high quality firm and will to choose to buy from the low quality firm with probability 1/2 if the other firm produces high quality and sets its equilibrium price. Accordingly, the payoff of deviating in this way is equal to $(1-\alpha)(c_H+\theta-c_L)/2$. This is smaller than the equilibrium payoff if $2(v_H-v_L)<c_H+\theta-c_L$. Given the condition in the proposition, this is the case for all values of θ.

(e) The low quality firm could set a price $p_L>c_H+\theta$. This is a situation similar to (c), the only difference being that the firm considers charging even higher prices. Following the argument under (c), it is clear that the consumer will not buy at such a high price. Accordingly, the low quality firm is not better off deviating from its equilibrium strategy.

Let us now consider the high-quality firm. In equilibrium, the expected payoff of this type of firm is given by $\alpha\theta/2$. We consider two possible deviations.

(a) A deviation to a price $p_H \leq c_H$ yields a non-positive payoff and is therefore not undertaken.

(b) For a deviation to a price p_H such that either $c_H N<p_H \leq c_H+\theta$ or $p_H>c_H+\theta$, we can give an identical argument as above under (c) and or (e): the consumer will not buy at this price and, hence, this deviation leads to a lower payoff.

We can conclude that it is not profitable for any type of firm to deviate from their equilibrium strategies. Q.E.D.

The mixed strategy distribution of the low quality firm has the same structure as the mixed strategies we have encountered before. This becomes apparent if we recall that earlier we have assumed that the cost of production is zero (set $c_L = 0$) and that the maximum price a firm can set is not anymore given by v, but by the price the firm can set to distinguish itself from a high quality firm, i.e., replace v by $c_H + \theta - (v_H - v_L)$. If we substitute these values into the expression of the mixed strategy equilibrium, we have the same expression as in Section 2.

Observe that the more the consumer believes that out-of-equilibrium prices are set by high quality firms (the smaller α'), the more competitive the industry and the lower the prices set in equilibrium. Following the same logic as above we can argue that if $\alpha' = 0$, a unique separating equilibrium exists in which the high quality firm set a price equal to its marginal cost.

(iii) Sequential Decisionmaking

In this last extension, N firms submit bids publicly in sequence from Firm 1 to Firm N, rather than simultaneously.

First, consider what happens in the Bertrand model with no uncertainty— the special case of $\alpha = 1$. There are two classes of equilibria.

In the first class of equilibrium, at least one of the first $N-1$ firms chooses $p = 0$, and consumers buy from firms charging $p = 0$. Profits are zero, and the outcome is the same as in the simultaneous Bertrand model.

In the second class of equilibria, the first $N-1$ firms choose prices in a set with minimum $p_{min} > 0$ and the last firm chooses $p_N = Min\{p_{min}, v\}$. The consumers all choose to buy from firm N. Profits are zero for all firms except Firm N, which has positive profit.

The second class of equilibrium is counter-intuitive. For concreteness, consider the particular member of the class in which all firms offer the price v. None of the first N firms have any incentive to deviate. If a firm deviates to $p = 0$, his profit remains zero. If a firm deviates to any price between 0 and v, Firm N will respond with the same price and capture the market, so the deviating firm's profit remains zero. Firm N clearly has no incentive to deviate. And the consumers have no incentive to deviate because all firms charge the same price.

This is, to be sure, a weak Nash equilibrium, which is why it is counterintuitive. No Nash equilibrium exists, however, in which consumers are not indifferent about where they buy, and in which more than one firm earns positive profits. This is the standard open-set problem; if consumers did not follow this behavior, the last firm would choose to undercut its lowest competitor by an infinitesimal amount and gain the entire market.[5] It seems reasonable, however, to prefer an equilibrium in which players behave symmetrically, when such an equilibrium exists, and with that assumption, the only equilibrium is the symmetric one in the first class with $p = 0$ for all firms, and consumers evenly divided among them.

Next, consider the Bertrand model in which each firm is active only with probability α. The last active firm will undercut any previous firm in the sequence that has offered $p > 0$. Thus, if any later firm in the sequence is active, the earlier firms' payoffs will all equal zero regardless of their bid. If no later firm is active, an early firm will win the market, and might as well bid $p = v$. Since if $\alpha < 1$ there is a positive probability that all later firms will be inactive, every firm bids $p = v$. Consumers buy from the last firm bidding, again, to resolve the open-set problem.

The properties of the equilibrium in the sequential model are somewhat bizarre. Even a tiny amount of uncertainty reduces a continuum of equilibria to a unique equilibrium. Moreover, the most plausible level of profits rises from zero to the monopoly level.

What this illustrates is the tremendous power of open-cry auctions in revealing information. When there is no uncertainty, this does not make much difference. When there is uncertainty about the number of firms, however, the open-cry auction resolves that uncertainty, giving the last bidder, in particular, a tremendous advantage. Earlier bidders know they cannot overcome that advantage, so their only hope is that no later bidders will be active.

The sequential Bertrand model is, of course, not a typical open-cry auction, because the sequence of bidding is predetermined and each firm only gets one bid. In the classic English auction, each bidder can bid as often as he wishes. In the present context, this would result in a winning bid of $p = 0$ if at least two firms are active, whatever the value of α may be.

The caveat "if at least two firms are active", however, is important. With probability $N(1-\alpha)^{N-1}\alpha$, only one firm is active and the winning bid will be $p = v$. The expected industry profit is therefore $N(1-\alpha)^{N-1}\alpha v$, exactly the same as the profit given by equation (8) in the simultaneous game! This is the same result found in McAfee & Macmillan (1987). From the point of view of the buyer, the English auction has the advantage of pitting bidders against each other head to head, but the disadvantage of letting a bidder know if he has no competition. As a result, the English auction has much greater risk, and a risk-averse buyer would prefer simultaneous bids.

[5]See p. 103 of Rasmusen (1994) for a discussion of the open-set problem. Note that the first edition of that book does not contain any discussion of it.

4. Comparing Bertrand and Cournot

The dichotomy between competition in quantity and competition in price has existed for over a hundred years. Cournot (1838) proposed a model in which N firms simultaneously choose quantities and let the market determine the price. Bertrand (1883) pointed out that entirely different conclusions result if the firms choose prices simultaneously instead. Even though price competition seems to yield a more reasonable model, the quantity model gives more accurate outcomes. In this section, we will show that this feature of the Bertrand model disappears when uncertainty about the presence of competitors is taken into account.

To compare Bertrand and Cournot oligopoly, we use a simple linear demand function:

$$p\left(\sum_{i=1}^{N} q_i\right) = a - b\sum_{i=1}^{N} q_i. \tag{25}$$

Let us define $q(p)$ as the demand facing a monopolist at a price of p, so

$$q(p) = \frac{a}{b} - \frac{p}{b}. \tag{26}$$

The monopoly price equals $a/2$ and quantity demanded is then $a/2b$,

What we will do in this section is to compute the expected profits from Cournot and Bertrand for different levels of N, to obtain some idea of the effects of concentration in each.

Bertrand equilibrium

As we have seen above, the industry profits in the Bertrand model with uncertainty are given by

$$\pi_{bertrand} = \frac{N\alpha(1-\alpha)^{N-1}\frac{a^2}{4b}}{1-(1-\alpha)^N}. \tag{27}$$

Thus, adding uncertainty eliminates the discontinuous behavior of the original Bertrand model. Profits are positive, but the expected price and profits decrease smoothly in the number of firms, as in the Cournot model that we will next analyze.

Cournot Equilibrium

Now let us compute the Cournot equilibrium. Let q^* be the Cournot output we are trying to determine. Then,

$$\begin{aligned} \pi_i(q_i) &= \alpha^{N-1}[p(q_i + (N-1)q^*)]q_i + (1-\alpha)\alpha^{N-2}[p(q_i + (N-2)q^*)]q_i \\ &+ (1-\alpha)^2\alpha^{N-3}[p(q_i + (N-3)q^*)]q_i + ... \\ &= \sum_{j=0}^{N-1}(1-\alpha)^j\alpha^{N-j}[p(q_i + (N-j)q^*)]q_i \end{aligned} \tag{28}$$

The first order condition is

$$\begin{aligned} \frac{d\pi_i(q_i)}{dq_i} &= \sum_{j=0}^{N-1}(1-\alpha)^j\alpha^{N-j}[-bq_i + a - b(q_i + (N-j)q^*)] = 0 \\ &= \sum_{j=0}^{N-1}(1-\alpha)^j\alpha^{N-j}[a - b(N-j+2)q^*] \\ &= \left(\sum_{j=0}^{N-1}(1-\alpha)^j\alpha^{N-j}\right)a - \left(\sum_{j=0}^{N-1}(1-\alpha)^j\alpha^{N-j}(N-j+2)\right)bq^*, \end{aligned} \tag{29}$$

so

$$q^* = \frac{\left(\sum_{j=0}^{N-1}(1-\alpha)^j\alpha^{N-j}\right)a}{\left(\sum_{j=0}^{N-1}(1-\alpha)^j\alpha^{N-j}(N-j+2)\right)b} \tag{30}$$

Note that if $\alpha = 1$, this boils down to $q^* = \frac{a}{(N+1)b}$. Adding incomplete information makes no great difference to the Cournot model. If some firms might not be active, each active firm produces somewhat more than it would have otherwise, but there is no qualitative shift in the equilibrium.

From the quantity we can get the expected profit conditional upon there being at least one firm in the market.

$$\pi_{Cournot} = \sum_{j=1}^{N+1}\left(\frac{(1-\alpha)^{N-j}\alpha^j}{1-(1-\alpha)^N}\right)p(jq^*)jq^* \tag{31}$$

Note that equation (31) is conditional upon $N * q^*$ not being so large as to drive the price to zero, which might rationally happen, since a firm would be willing to accept a price of zero occasionally as the result of all N firms coincidentally being active and producing a large amount.

Comparisons

Profits are positive but fall with the number of firms in both the Bertrand and Cournot equilibria. The question is how fast profits fall. This is best seen with a numerical example. Let $a = 110$, $b = 1$, and $N = [0, 7]$ for $\alpha = 1$ and $\alpha = .9$. Table 1 and Figures 2 through 4 show the levels of profits. Table 1 consists of industry profits under the Bertrand and Cournot models with certainty and with $\alpha = .9$, and its numbers are repeated graphically in Figure 2. Figures 3 and 4 show how the profit/concentration relationship changes for different values of α in the two models.[6]

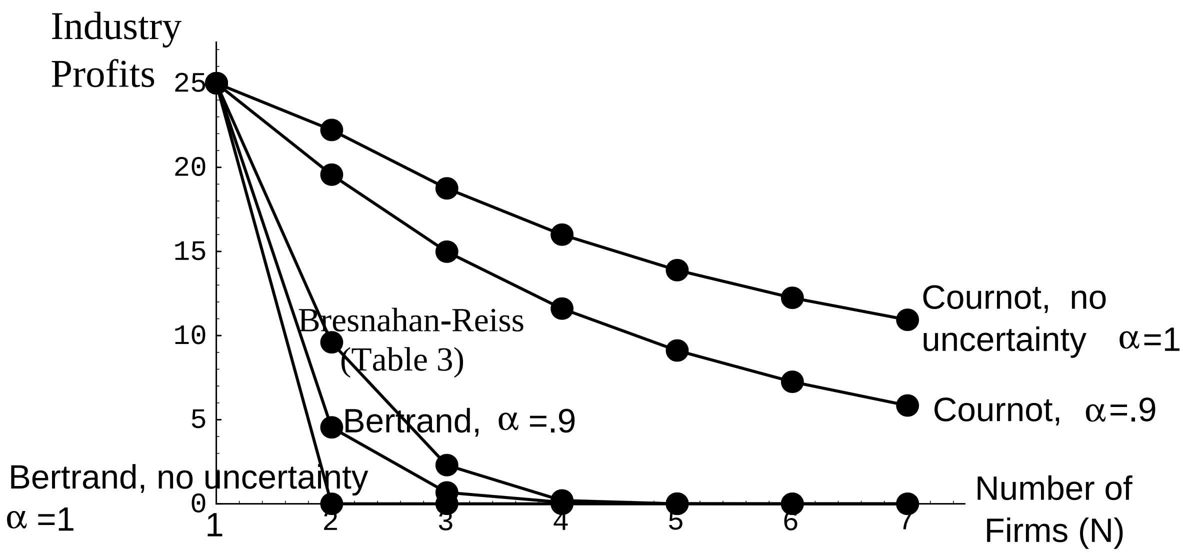

Figure 2: Bertrand and Cournot Profits

Consider first the Cournot model. Table 1 and Figure 2 show that a small amount of uncertainty makes little difference in the Cournot model, though, oddly enough, industry profits actually fall when the expected percentage of active firms declines. Under Cournot competition, a firm expands its output when it expects fewer rivals to be helping push down the price. Uncertainty over the number of rivals ends up increasing average output and reducing profits, a peculiar result. Figure 3 shows that this is a very delicate conclusion. For $N > 3$, profits rise when α falls from 0.4 to 0.1, but also rise when α rises from .4 to .7, but the comparative statics can switch if N is smaller. Conflicting forces are at work in Cournot equilibrium, and the net result is sensitive to the particular assumptions of the model.[7]

[6]In every case, expected industry profits are conditional upon at least one firm being active. When $\alpha = 0$, this is to be interpreted as the probability zero (but possible) event that one firm is active and the expected number of other firms is zero.

[7]The result is reminiscient of the peculiarities of profit per firm in the Cournot model, which can (but not always) give rise to an incentive for a Cournot firm to split in two to increase its profits. See Salant, Switzer, and Reynolds (1983) for that effect.

Number of Firms N	1	2	3	4	5	6	7
Bertrand, $\alpha = 1$	25.0	0.0	0.0	0.0	0.0	0.0	0.0
Bertrand, $\alpha = .9$ (eq. (27))	25.0	5.6	0.9	0.1	0.02	0.003	0.0003
Cournot, $\alpha = 1$	25.0	22.2	18.8	16.0	13.9	12.2	10.9
Cournot, $\alpha = .9$ (eq. (31))	25.0	19.6	15.0	11.6	9.1	7.3	5.8

Table 1: Industry Profits for Different Concentration Levels[8]

Uncertainty is much more important in the Bertrand model, and the comparative statics are more consistent and intuitive. Table 1 and Figure 2 show that a small amount of uncertainty changes the Bertrand model in a small but crucial way, because profits do become positive and monotonic in the number of firms. The sharp fall in profits moving from monopoly to duopoly under certainty is not so unreasonable as it looks. It is extreme, but it is a limiting result as α goes toone, as Figure 4 illustrates.

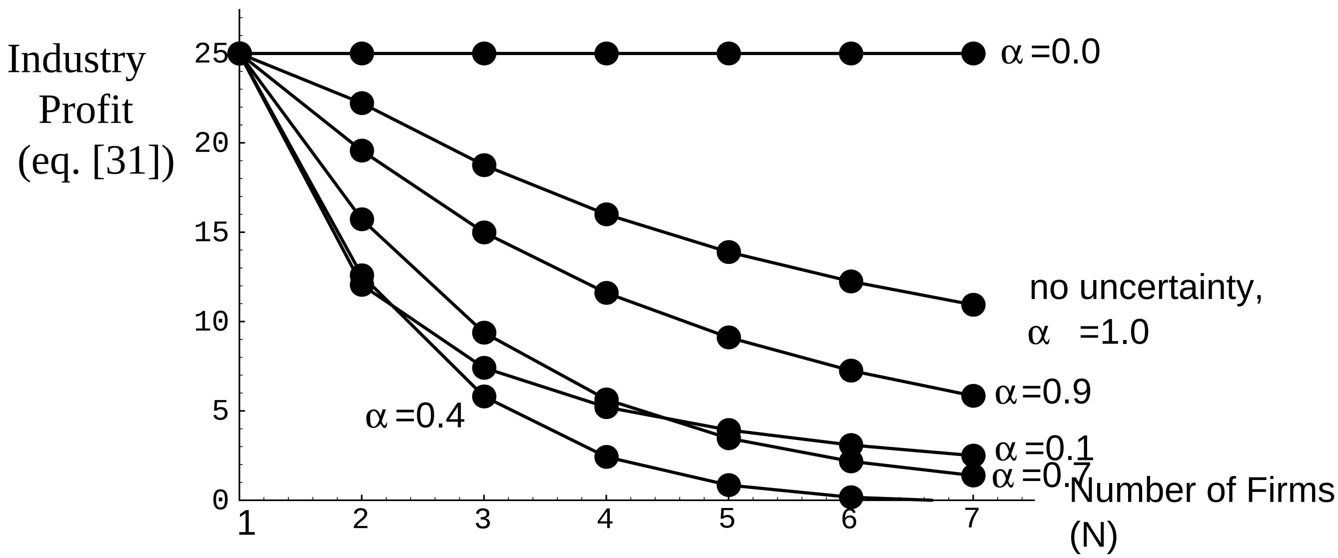

Figure 3: Cournot Profits For Different Probabilities of Activity α and Numbers of Firms N
(conditional on at least one firm being active)

[8]Numerical calculations and figure-drawing used *Mathematica*. Values are rounded.

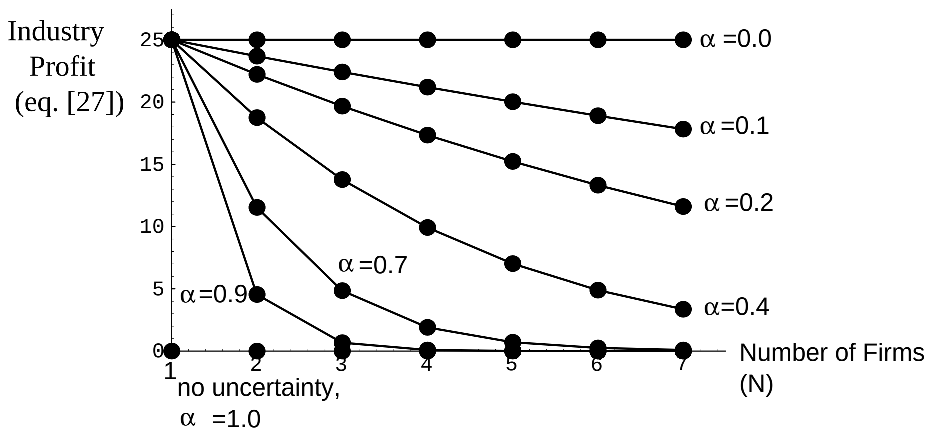

Figure 4: Bertrand Profits For Different Probabilities of Activity α and Numbers of Firms N
(conditional on at least one firm being active)

Let us also consider the shape of the profit-concentration paths. All the curves in Figures 2 through 4 have convex shapes, if only weakly in the limiting cases, but the curvatures, and therefore the empirical implications, are different. As Figure 4 and Table 1, in particular, show, profits decline much more rapidly in Bertrand than in Cournot. For the parameters chosen, industry profits fall from the monopoly level of 25 to duopoly profits of 5.6, triopoly profits of 0.9, and negligible levels thereafter. Cournot profits show a much more uniform decline as concentration falls.

Comparison of Figures 3 and 4 shows that for smaller values of the activity probability α the Bertrand profit path becomes flatter and the Cournot path, perhaps more curved, but even at extreme values Cournot does not generate such sharp differences from the addition of one firm to the market.

For most modelling purposes, these models are building blocks, and such subtle differences in the profit-concentration path are unimportant. They are interesting, however, if one wishes to consider Bertrand and Cournot as serious oligopoly models in their own right. Empirically, then, how do profits react to the number of firms? Do they decline to zero with duopoly and then stay constant, as in the original Bertrand model? Do they decline smoothly, as either version of the Cournot model would suggest? Or do they decline rapidly, as the Bertrand model with uncertainty would suggest?

Measuring the relationship between profits and concentration is an old exercise now in some disrepute.[9] The difficulty is that the usual unit of observation has been the industry. This is natural enough, since one needs a measurement of concentration for each observation. Comparing accounting profits across industries is fraught with danger, however, since accounting profits differ from economic profits in ways that depend on the industry chosen and which are very likely to be correlated with technology, and hence with concentration. Moreover, it is not clear that the concentration-profits path is even the same across industries.

A clever recent approach to the same problem is that of Bresnahan & Reiss (1991). They took the unit of observation to be the market for a particular product in a particular small town, rather than for many products over the entire United States, and they looked at market size rather than directly at profits. They collected data on the size of a town and the number of dentists there, for example. If a town is very small—say, 500 people—- it will have no dentist, since a dentist incurs a fixed cost and could not make any profit there even as a monopoly. If it grows to 800 people, it will have one dentist, since the profits are enough for monopoly, but entry by a second dentist would drive them negative. If the town grows to 1,600 people, however, it may still have only one dentist— if entry by the second dentist would not just split the industry profits, but reduce them.

Number of Firms N	1	2	3	4	5
Doctors	0.88	1.75	1.93	1.93	1.83
Dentists	0.71	1.27	1.39	1.36	1.28
Druggists	0.53	1.06	1.68	1.92	1.88
Plumbers	1.43	1.51	1.51	1.55	1.49
Tire Dealers	0.49	0.89	1.14	1.19	1.22

Table 2: Bresnahan-Reiss Entry Thresholds s_i: Original (1,000's of inhabitants)[10]

Bresnahan and Reiss used this approach to estimate the thresholds for entry in

[9] See pp. 349-366 of Carlton & Perloff 's 1994 industrial organization text for a good discussion of the problems of the profits-concentration literature.

[10] Calculated from Table 5A of Bresnahan & Reiss (1991). Note that the entry of .79 in the second row of their original paper is a mistake and should be 1.09, and their Figure 4 illustrates s_i/s_5, not the s_5/s_i in the legend.

small markets for a number of industries. Table 2 shows these thresholds in thousands of inhabitants per firm. Table 3 rescales the same numbers to be very roughly comparable with the numerical example used earlier in this paper.[11] The rescaling is somewhat arbitrary, since the theory of Bresnahan and Reiss is that some quasi-rents remain to cover fixed cost even when the minimum scale for entry flattens out, but it creates a comparison measure for how the intensity of competition changes with the number of firms.

Number of Firms N	1	2	3	4	5
Doctors	25.0	4.3	0.0	0.0	0.0
Dentists	25.0	4.4	0.0	0.0	0.0
Druggists	25.0	15.5	4.3	0.0	0.0
Plumbers	25.0	8.3	8.3	0.0	0.0
Tire Dealers	25.0	11.3	2.7	1.0	0.0
Average	25.0	9.6	2.3	0.2	0.0

Table 3: Bresnahan-Reiss Entry Thresholds: Rescaled $(\frac{25(s_m - s_i)}{(s_m - s_1)})$

What is significant is how profits flatten out, even though the choice of 0 as the flat level in Table 3 is arbitrary.[12] The empirical result that going from one firm to two is much more important than going from two to three, and that full-fledged competition kicks in very quickly matches the Bertrand model with uncertainty very well, and is inconsistent with the Cournot model.

6. Concluding Remarks

The Bertrand model with uncertainty about the number of competitors is simple, but its properties are both interesting and useful. the extreme transition from monopoly to competition found in the standard Bertrand model disappears. Expected profits are positive, but decline with the number of firms in the industry, and decline in a way that empirical work suggest is more realistic than in the Cournot model. We have tried to show that the model is useful both as a simple decription of oligopoly and as a building block for other topics in industrial organization, as in Gwin (1997), and Janssen and Van Reeven 1998).

[11] Table 3's rescaling uses the following procedure. Define the monopoly level of profits in an industry to be 25, and the competitive level to be 0. Assume that when s_i reaches its maximum level s_m over $[1, 5]$, the competitive level of profits is reached and any further changes are measurement error. Apply the conversion formula $s_i^* = \frac{25(s_m - s_i)}{(s_m - s_1)}$, and Table 3 results.

[12] Add 9.1 to each entry in Table 3, and the profit at $N = 5$ is 9.1, as with Cournot competition and $\alpha = 0.9$ in Table 1, but the shape is still more like that of Bertrand competition.

Appendix on Convexity

This appendix repeats some of the analysis of Section 2 but includes several more steps of algebra.

Let π_a denote expected industry profit given that at least one firm is active. The profit in equation (8) can be written as

$$\sum_{i=1}^{N} \pi_i = N\alpha(1-\alpha)^{N-1}v = (1-\alpha)^N(0) + [1-(1-\alpha)^N]\pi_a, \tag{32}$$

yielding

$$\pi_a = \frac{N\alpha(1-\alpha)^{N-1}v}{1-(1-\alpha)^N}. \tag{33}$$

To see how industry profit changes with N, note that after some algebra,

$$\frac{d\pi_a}{dN} = \left[\frac{(1-(1-\alpha)^N) + Nlog(1-\alpha)}{(1-(1-\alpha)^N)^2}\right]\left[\alpha(1-\alpha)^{N-1}v\right] \tag{34}$$

a derivative which is well-defined even though only integer values of N have an economic interpretation. The sign of expression (11) is the sign of

$$1-(1-\alpha)^N + Nlog(1-\alpha). \tag{35}$$

For $N=1$, expression (35) becomes $\alpha + log(1-\alpha)$, which is negative because $\alpha < 1$. For larger N, expression (35) becomes even more negative, because its derivative is $-(1-\alpha)^N log(1-\alpha) + log(1-\alpha) = log(1-\alpha)[1-(1-\alpha)^N] < 0$. Thus,

$$\frac{d\pi_a}{dN} < 0, \tag{36}$$

and profits fall as the number of firms increases.

We can say more. The first derivative from (11) can be rewritten as

$$\frac{d\pi_a}{dN} = \alpha v \left\{\frac{(1-\alpha)^{N-1}}{1-(1-\alpha)^N} + \frac{Nlog(1-\alpha)}{[1-(1-\alpha)^N]^2(1-\alpha)^{N-1}}\right\}. \tag{37}$$

The derivative of this is

$$\begin{aligned}\frac{d^2\pi_a}{dN^2} &= \alpha v \left\{\frac{[1-(1-\alpha)^N](1-\alpha)^{N-1}log(1-\alpha)+(1-\alpha)^N(1-\alpha)^{N-1}log(1-\alpha)}{[1-(1-\alpha)^N]^2} + \right.\\ &\left.\frac{[(1-\alpha)^{N-1}log(1-\alpha)+(1-\alpha)^{N-1}Nlog^2(1-\alpha)][1-(1-\alpha)^N]^2+2[1-(1-\alpha)^N]/\{(1-\alpha)^N log(1-\alpha)[N(1-\alpha)^{N-1}log(1-\alpha)]\}}{[1-(1-\alpha)^N]^4}\right\}.\\ &= \frac{(1-\alpha)^{N-1}log(1-\alpha)}{[1-(1-\alpha)^N]^2} + \frac{(1-\alpha)^{N-1}Nlog^2(1-\alpha)[1-2(1-\alpha)^N]^2+(1-\alpha)^{2N}+2(1-\alpha)^N-2(1-\alpha)^{2N}}{[1-(1-\alpha)^N]^4}.\\ &= \frac{(1-\alpha)^{N-1}log(1-\alpha)}{[1-(1-\alpha)^N]^2} + \frac{(1-\alpha)^{N-1}Nlog^2(1-\alpha)[1-(1-\alpha)^{2N}]}{[1-(1-\alpha)^N]^4}.\\ &= \frac{(1-\alpha)^{N-1}log(1-\alpha)}{[1-(1-\alpha)^N]^2}\left\{1 + \frac{Nlog(1-\alpha)[1+(1-\alpha)^N]}{1-(1-\alpha)^N}\right\}.\end{aligned} \tag{38}$$

The first term of this last expression is negative because $log(1-\alpha)$ is negative.

The second term has the same sign as

$$1-(1-\alpha)^N + Nlog(1-\alpha)[1+(1-\alpha)^N. \tag{39}$$

We already found that for all N, expression (12) is negative, i.e.,

$$1-(1-\alpha)^N + Nlog(1-\alpha) < 0. \tag{40}$$

The third term of expression (39) is negative, and $[1+(1-\alpha)^N > 1$, so $Nlog(1-\alpha)[1+(1-\alpha)^N$ is more negative than $Nlog(1-\alpha)$, so expression (39) must be negative if inequality (40) is true. Hence, since both terms of the last expression in (38) are negative, their product must be positive, and we can conclude that

$$\frac{d^2\pi_a}{dN^2} > 0. \tag{41}$$

This means that profits are convexly decreasing in the number of firms in the industry, so the shape shown in the numerical examples graphed in Figure 1 would be found for any example.

References

Baye, Michael, Dan Kovenock, & Casper de Vries (1997) "A General Linear Model of Contests," working paper, Indiana University, Department of Business Economics and Public Policy, Http://php.indiana.edu/~mbaye/linear.pdf.

Baye, Michael & John Morgan (1997a) "Information Transmission, Information Acquisition, and Price Dispersion in 'Thin' Markets," Working paper, Indiana University, Kelley School, Dept. of Business Economics and Public Policy.

Baye, Michael & John Morgan (1997b) "Necessary and Sufficient Conditions for Existence and Uniqueness of Bertrand Paradox Outcomes, " Working paper, Indiana University, Kelley School, Dept. of Business Economics and Public Policy. Http://php.indiana.edu/~mbaye/bert.pdf.

Bertrand, Joseph (1883) "Rechercher sur la Theorie Mathematique de la Richesse." *Journal des Savants.* September 1883. 48: 499-508.

Bresnahan, Timothy & Peter Reiss (1991) "Entry and Competition in Concentrated Markets." *Journal of Political Economy.* October 1991. 99: 977-1009.

Burdett, Kenneth & Kenneth Judd (1983) "Equilibrium Price Dispersion," *Econometrica*, July 1983, 51: 981-996.

Carlton, Dennis & Jeff Perloff (1994) *Modern Industrial Organization*, second edition, New York: HarperCollins College Publishers, 1994.

Cho, In-Koo & David Kreps (1987) "Signaling Games and Stable Equilibria" *Quarterly Journal of Economics.* May 1987. 102: 179-221.

Cournot, Augustin (1838) *Recherches sur les Principes Mathematiques de la Theorie des Richesses.* Paris: M. Riviere & C., 1938. Translated in *Researches into the Mathematical Principles of Wealth.* New York: A.M. Kelly, 1960.

Edgeworth, Francis (1897) "La teoria pura del monopolio" *Giornale degli economisti.* 1925. 40: 13-31. Translated in Edgeworth, Francis, *Papers Relating to Political Economy*, Vol. I. London: Macmillan, 1925, 111-42.

Gwin, Carl (1997) "Special Pricing Agreements as Vertical Restraints," in *Three Common Business Practices as Reponses to Asymmetric Information Problems*, Ph.D. Dissertation, Indiana University School of Business, July 1997.

Janssen, Maarten C.W. and Peran Van Reeven (1998) "Market Prices and Illegal Practices." *International Review of Law and Economics* (forthcoming).

Kreps, David and Jose Scheinkman (1983) "Quantity Precommitment And Bertrand Competition Yield Cournot Outcomes." *Bell Journal of Economics.* 14: 326-337.

Lacasse, Chantale (1995) "Bid Rigging and the Threat of Government Prosecution," *RAND Journal of Economics*, Autumn 1995, 26: 398-417.

Lang, Kevin, and Robert Rosenthal (1991), " The Contractors' Game" *RAND Journal of Economics*, Autumn 1991, 22: 329-338.

Levitan, Richard & Martin Shubik (1972) "Price Duopoly and Capacity Constraints." *International Economic Review*. February 1972. 13: 111-22.

McAfee, R. Preston & John Macmillan (1987) "Auctions with a Stochastic Number of Bidders," *Journal of Economic Theory*, 43: 1-19.

Rasmusen, Eric (1994) *Games and Information*. 2nd edition. Oxford: Basil Blackwell, 1994.

Salant, Stephen, Sheldon Switzer & Robert Reynolds (1983) "Losses from Horizontal Merger: The Effects of an Exogenous Change in Industry Structure on Cournot-Nash Equilibrium," *Quarterly Journal of Economics*. May 1983. 98: 185-99.

Spulber, Daniel (1995) "Bertrand Competition When Rivals' Costs are Unknown," *Journal of Industrial Economics*. March 1995. 43: 1-11.