Economics
Discussion Paper Series
EDP-0917

On the existence of Bayesian Bertrand equilibrium

Robert R. Routledge

September 2009

Economics
School of Social Sciences
The University of Manchester
Manchester M13 9PL

On the existence of Bayesian Bertrand equilibrium

Robert R. Routledge*

Economics

School of Social Sciences

University of Manchester

robert.routledge@postgrad.manchester.ac.uk

September, 2009

Abstract

We analyze the classical model of Bertrand competition in a homogeneous good market with constant marginal costs and uncertainty regarding rivals' costs. First, we show that there exists a mixed strategy Nash equilibrium under the conventional equal sharing rule. Second, we illustrate the result for the case of piecewise-affine market demand.

JEL Classification: D43, L13

Keywords: Bertrand competition; cost uncertainty; mixed strategies

*Correspondence address: Economic Studies, Arthur Lewis Building, University of Manchester, Oxford Road, United Kingdom, M13 9PL. I should like to thank Chris Birchenhall and Alejandro Saporiti for their helpful comments. All remaining errors are the responsibility of the author.

1 Introduction

Up until the last twenty years the Bertrand model of price competition was associated with the striking result that if two firms compete by simultaneously setting prices, both firms have symmetric constant returns to scale cost functions, and the market demand possesses a finite choke-off price, then the unique Nash equilibrium is that both firms price at marginal cost and earn zero profits. This, of course, is the famous 'Bertrand paradox'.[1] However, subsequent research has revealed that the Bertrand model of price competition has other notable results. Dastidar (1995) showed that if firms have strictly convex costs, consistent with decreasing returns production technology, then there exists a continuum of pure strategy price equilibria in which the firms earn strictly positive profits. This result also holds if firms play mixed strategies. Hoernig (2002) has shown then when firms have strictly convex costs then there exist different types of mixed strategies with firms placing probability mass on a finite number of prices, or playing mixed strategies with continuous supports.

When firms have increasing returns to scale cost functions, the literature indicates that the existence of pure strategy Bertrand equilibrium is problematic.[2] Dastidar (2006) has shown that when firms have increasing returns to scale cost functions then there does not exist any pure strategy price equilibrium under the conventional equal sharing rule. However, it is well-known that if the sharing rule is altered so that a single firm is selected randomly from the set of firms tieing at the minimum price to serve all the market demand, the so-called "winner-takes-all" sharing rule, then this restores the existence of pure strategy equilibrium (Vives, 1999, p.119). The existence of a mixed strategy equilibrium with increasing returns to scale costs has, up until recently, been an open question. In their survey of the literature Baye and Kovenock (2008) proved that there

[1] If no finite choke-off price exists, and monopoly revenues are unbounded, then in addition to the pure strategy zero profit Nash equilibrium, there exist a continuum of atomless mixed strategy equilibria in which the firms earn positive profits (Kaplan and Wettstein, 2000).

[2] As Bertrand competition is a game with discontinuous payoffs, the Glicksberg-Fan theorem cannot be used to establish existence, and more recent existence results for discontinuous games, such as Dasgupta and Maskin (1986) and Reny (1999), are not applicable because payoffs are neither quasi concave, sum upper semicontinuous nor reciprocally upper semicontinuous.

does not exist a mixed strategy equilibrium either.

An extension of the Bertrand model is to consider the case when firms have asymmetric costs. If both firms have constant marginal costs, and the high marginal cost is below the profit maximizing price for the low-cost firm, it has recently been shown that the price equilibrium has the low-cost firm setting price equal to the higher marginal cost, and the high-cost firm randomizing uniformly on an interval above (Blume, 2003).[3] Related to asymmetric costs is research which analyzes the Bertrand model when costs are uncertain. Typically, it is assumed each firm knows their own cost type but does not know the cost type of their rivals. Spulber (1995) showed that if the cost function is parameterized, and the parameter drawn from a continuous probability distribution, then if firms are uncertain about their rivals' cost profiles, they price above marginal cost and make positive profits. Vives (1999, p.230) analyzed the general case where firms have constant marginal costs drawn from a continuous probability distribution with compact support, and showed by the log-supermodularity of the expected profit function that there always exists a Bayesian Nash equilibrium in which firms set prices contingent upon cost type.

However, there is a notable gap in the research. There are no equilibrium existence results for the classical Bertrand model when there is discrete cost uncertainty. A survey of textbook questions on market competition reveals that the problem of finding a Bayesian Nash equilibrium in the homogeneous good Cournot model with discrete cost uncertainty is a frequently encountered question (Mas-Colell et al., 1995, p.265).[4] Bertrand competition with differentiated goods and discrete cost uncertainty is also a commonly encountered textbook problem (Tirole, 1988, p.362). But the case of homogeneous good Bertrand price competition, with constant returns to scale production technology, and discrete cost uncertainty, does not appear in any textbooks or in the professional literature. The aim here is to fill this gap. The main result (Proposition 2) is that there exists a mixed strategy equilibrium. The next section of the paper presents the model

[3]For a long time this simple case of asymmetric costs did not have a well-defined solution. It was typical to assume the price space was discrete and the low-cost firm would price at some small unit, a minimum currency, below the high marginal cost. See, for example, Tirole (1988, p.211).

[4]Although Einy et al. (2009) have shown that when firms have differential information regarding costs and/or demand, and there are a finite number of states, then a pure strategy Bayesian Cournot equilibrium may fail to exist.

and result. The final section draws some conclusions regarding future research.

2 The Model

Consider the market for a homogeneous good in which there are $N = \{1, ..., n\}$, $n \geq 2$, firms which compete by simultaneously and independently setting prices. The market demand $D : \Re_+ \rightarrow \Re_+$ is C^2 and $\exists$ positive finite real numbers P^{Max}, Q^{Max}, satisfying $D(P^{Max}) = 0$ and $D(0) = Q^{Max}$. Also, $D'(P) < 0 \; \forall P \in (0, P^{Max})$. Firms' cost functions are derived from constant returns to scale production technology and take one of two forms: high or low. That is, $C_i(Q) = c_i Q$ with $i \in \{H, L\}$ and $0 < c_L < c_H < P^{Max}$. Each firm has probability θ of having marginal cost c_L and probability $1 - \theta$ of having marginal cost c_H with $\theta \in (0, 1)$. It is assumed each firm knows their own cost type but does not know the cost type of their rivals. Firms supply all the demand they face and maximize their expected profits.[5] Ties at the lowest price are resolved by firms sharing the market demand equally. Define $\pi_i(P)$ to be the monopoly profit of a firm with cost type i, and $\hat{\pi}_i(P, m)$ to be the shared profit of a firm with cost type i when it ties with $m - 1$ firms at the lowest price:

$$\pi_i(P) = (P - c_i)D(P) \tag{1}$$

$$\hat{\pi}_i(P, m) = \frac{1}{m}(P - c_i)D(P) \tag{2}$$

Finally, we assume that $\pi_i(P)$ and $\hat{\pi}_i(P, m)$ are strictly concave in price. All the assumptions employed here are standard specifications of the Bertrand model apart from the uncertainty regarding costs. In the classical Bertrand model the firms have the same cost type and the price-setting game is a one of complete information. Here, each firm only knows their own cost type and the probability distribution over the possible cost types of their rivals. As a result, the price setting game becomes a game of incomplete information. Let $\mathbb{P} = [0, P^{Max}]$ denote the pure strategy price space, and let $\mathcal{G}$ denote the game in which firms simultaneously and independently set prices. In order to simplify

[5] This is what distinguishes Bertrand competition from Bertrand-Edgeworth competition. In Bertrand-Edgeworth competition firms may refuse to supply all the demand forthcoming at any price. For a succinct summary, see Vives (1999, Ch.5).

the main result define the following:

$$P^* = \arg\max_{P \in \mathbb{P}} \pi_L(P)$$

$$\hat{P} = \min\{c_H, P^*\}$$

In words, P^* is the monopoly price for the low-cost firm, and $\hat{P}$ is the minimum of the high marginal cost and P^*.

Lemma 1 *There exists a unique* $\tilde{P} \in (c_L, \hat{P})$ *s.t.* $\pi_L(\tilde{P}) = (1-\theta)^{n-1}\pi_L(\hat{P})$.

Proof. First, note that $0 < (1-\theta)^{n-1}\pi_L(\hat{P}) < \pi_L(\hat{P})$. As the profit function is continuous, and $\pi_L(c_L) = 0$, the intermediate value theorem[6] guarantees $\exists$ a $\tilde{P} \in (c_L, \hat{P})$ such that $\pi_L(\tilde{P}) = (1-\theta)^{n-1}\pi_L(\hat{P})$. The strict concavity of the profit function implies $\pi'_L(P) > 0$ for all $P \in (c_L, \hat{P})$. This ensures that $\tilde{P}$ is unique. ■

Proposition 1 $\mathcal{G}$ *does not possess a pure strategy Bayesian Nash equilibrium.*

Proof. Start by assuming there does exist a pure strategy equilibrium. There are two cases to consider. The first case is when the firms play a symmetric pure strategy. The second case is when the firms play an asymmetric pure strategy.

Case 1: The firms play a symmetric pure strategy with the low-cost firms setting a price $P^L > c_L$ and the high-cost firms setting a price $P^H > c_H$, and suppose, without loss of generality, that $P^L < P^H$. That is, when low-cost firms are in the market, high-cost firms do not receive any demand. The expected payoff to the high-cost firm is $(1-\theta)^{n-1}\hat{\pi}_H(P^H, n)$. As $\pi_H(P) > \hat{\pi}_H(P, n)$ for all $P \in (c_H, P^{Max})$ one firm could deviate to a price $P^H - \epsilon > c_H$ and increase their expected profit. The possibility of undercutting means that in any equilibrium the high-cost firm must price at marginal cost: $P^H = c_H$. Now consider the low-cost firm. The expected profit is:

$$(1-\theta)^{n-1}\pi_L(P^L) + \sum_{r=1}^{n-1} \binom{n-1}{r} \theta^r (1-\theta)^{n-1-r} \hat{\pi}_L(P^L, r+1)$$

Which is strictly less than $\pi_L(P^L)$. This means one low-cost firm could deviate to a

[6]See, for example, Rudin (1976, p.93) Theorem 4.23.

price $P^L - \epsilon > c_L$ and increase its expected profit. As with the the high-cost firm, the possibility of undercutting leads to the consideration of $P^L = c_L$ as a possible equilibrium. However, this cannot be an equilibrium because if $P^L = c_L$ then one firm could deviate to a price $P \in (c_L, c_H)$ and earn expected profit of $(1-\theta)^{n-1}\pi_L(P) > 0$. Hence, there does not exist a symmetric pure strategy equilibrium.

Case 2: The firms play asymmetric pure strategies. The same reasoning as above means that the high-cost strategy must have at least two firms pricing at marginal cost, and all other firms either pricing at marginal cost or a higher price.[7] However, as no low-cost firms want to tie at the same price suppose, without loss of generality, that the pricing strategies of the low-cost firms are $c_L \leq P_1^L < P_2^L ... < P_n^L < c_H$. Consider the firms setting the lowest two prices. If $P_2^L \leq P^*$ then firm one could increase its expected profit by setting a price arbitrarily close to P_2^L. Assume that $P_1^L = P_2^L - \epsilon$. Then firm two would be able to increase its expected profit by pricing slightly below firm one. Similarly, if $P_2^L > P^*$, then firm one is best choosing $P_1^L = P^*$. Then firm two would be able to increase its expected profit by slightly undercutting firm one. This possibility of undercutting means that if there is an equilibrium then it must be symmetric with $P_i^L = c_L$ for $i = 1, ..., n$. As was shown in Case 1, pricing at marginal cost is not an equilibrium for the low-cost firm because one firm could deviate to a price $P \in (c_L, c_H)$ and earn expected profit of $(1-\theta)^{n-1}\pi_L(P) > 0$. Hence, there does not exist a pure strategy price equilibrium. ■

Proposition 2 *$\mathcal{G}$ possesses a mixed strategy Bayesian Nash equilibrium.*

Proof. We shall show that there exists an equilibrium with the low-cost firm playing an atomless mixed strategy on $[\tilde{P}, \hat{P}]$ and the high-cost firm pricing at marginal cost. In equilibrium the expected payoff to the low-cost firm is strictly positive and the expected payoff to the high-cost firm is zero. First, we show what the strategies of the firms are. Second, we show that no firm wants to deviate from these strategies.

Step 1: Let $F(P)$ be the distribution function which describes the mixed strategy the low-cost firm plays over $[\tilde{P}, \hat{P}]$. Let $\alpha \in \Re$ denote the expected payoff to the low-cost firm. As the low-cost firm must be indifferent between playing any price in the support

[7] This is the standard Bertrand equilibrium when there are two or more firms in the market.

of $F(P)$ we require:

$$\sum_{r=0}^{n-1} \binom{n-1}{r} \theta^r (1-\theta)^{n-1-r} (1-F(P))^r \pi_L(P) = \alpha \tag{3}$$

We can now check that for any $P \in [\tilde{P}, \hat{P}]$ there is an implied value of $F(P)$ which satisfies (3) and possesses all the required properties of a distribution function. First, we require that $F(\hat{P}) = 1$. Substituting this into (3) then gives $\alpha = (1-\theta)^{n-1}\pi_L(\hat{P}) > 0$. Equation (3) can be rewritten as:

$$\sum_{r=0}^{n-1} \binom{n-1}{r} \theta^r (1-\theta)^{n-1-r} (1-F(P))^r = \frac{(1-\theta)^{n-1}\pi_L(\hat{P})}{\pi_L(P)} \tag{4}$$

Evaluating the R.H.S of (4) at $\tilde{P}$, and using the result in Lemma 1, implies that $F(\tilde{P}) = 0$ on the L.H.S. Moreover, as the R.H.S is continuous and strictly decreasing in price, this implies that $F'(P) > 0$ for all $P \in [\tilde{P}, \hat{P})$. Hence, we know that there exists an atomless mixed strategy, described by $F(P)$, which gives the low-cost firm an expected payoff of $\alpha > 0$.[8]

Step 2: Now consider whether any firm can profitably deviate from playing $F(P)$ when low-cost, and pricing at marginal cost when high-cost. Consider the low-cost firm. Assume that $\hat{P} = c_H$. There are two possible deviations. The low-cost firm could deviate to a price $P' \in [0, \tilde{P})$. From Lemma 1 we know that $\pi_L(\tilde{P}) = \alpha$. The strict concavity of the profit function implies $\pi_L(P') < \alpha$. Second, the low-cost firm could deviate to a price $P' > c_H$. Then the expected profit is zero. Assume that $\hat{P} = P^*$. Then in addition to the possible deviations already covered the low-cost firm could deviate to a price $P' \in (P^*, c_H)$ and earn expected profit of $(1-\theta)^{n-1}\pi_L(P')$. The expected payoff from playing $F(P)$ is $\alpha = (1-\theta)^{n-1}\pi_L(P^*)$. As P^* is the profit maximizing price for the low-cost firm this is not a profitable deviation. Hence, the low-cost firm has no better strategy than to play $F(P)$. Now consider the high-cost firm. There are two possible price deviation. First, a high-cost firm could set a price $P' < c_H$. As the high-cost firm would then be serving any demand forthcoming at a price less than marginal cost, the expected profit would be less than zero. Second, a high-cost firm could set a price $P' > c_H$. Then

[8] Note that if $\hat{P} = c_H$ then payoff indifference holds almost everywhere in the support. I am grateful to Andreas Blume for clarifying this point.

this firm would be undercut with certainty and make zero profit, which is no better than pricing at marginal cost. Therefore we conclude that no firm can profitably deviate from playing $F(P)$ when low-cost and pricing at marginal cost when high cost. ■

We illustrate the result in the following numerical example:

Example *Consider a duopoly,* $n = 2$*, with piecewise-affine market demand* $D(P) = \max\{0, 10 - P\}$*. Each firm has either a high or low marginal cost with* $c_L = 1$ *and* $c_H = 3$*. The probability of having low-marginal cost is* $\theta = \frac{3}{7}$ *and the probability of having high-marginal cost is* $1 - \theta = \frac{4}{7}$*. Routine calculations reveal that* $P^* = 5\frac{1}{2}$*,* $\hat{P} = 3$*,* $\tilde{P} = 2$ *and* $\alpha = 8$*. The distribution function which describes the mixed strategy the low-cost firm plays satisfies:*

$$\frac{3}{7}(1 - F(P))\pi_L(P) + \frac{4}{7}\pi_L(P) = 8$$

This can be solved for the distribution function:

$$F(P) = \frac{7}{3}\left[1 - \frac{8}{(P-1)(10-P)}\right]$$

It can be checked that $F(2) = 0$*,* $F(3) = 1$ *and* $F'(P) > 0$ *for all* $P \in [2, 3)$*. Hence we know from Proposition 2 that the low-cost firm playing* $F(P)$ *on* $[2, 3]$ *and the high-cost firm pricing at marginal cost is a Bayesian Nash equilibrium.*

3 Conclusion

In this paper we have shown that the classical Bertrand model possesses a mixed strategy Bayesian Nash equilibrium when there is discrete cost uncertainty. Previous research showed that the Bertrand model possesses mixed strategy equilibria when monopoly revenues are unbounded (Kaplan and Wettstein, 2000), or when costs are strictly convex (Hoernig, 2002), but as far as the author is aware, this is the first paper, with the exception of Blume (2003), to show the existence of a mixed strategy equilibrium with bounded demand and constant marginal costs. Future research should explore the existence of equilibrium with incomplete information when firms compete in different ways. For example, Weibull (2006) has examined the existence of equilibrium in repeated price competition when firms have convex costs. It would be of interest to extend these results to consider repeated competition with incomplete information.

References

[1] Baye, M. and D. Kovenock, 2008, Bertrand competition, The New Palgrave Dictionary of Economics.

[2] Blume, A., 2003, Bertrand without fudge, Economics Letters, 78, 167-8.

[3] Dasgupta, P. and E. Maskin, 1986, The existence of equilibrium in discontinuous economic games, Review of Economic Studies, 53, 1-26.

[4] Dastidar, K., 1995, On the existence of pure strategy Bertrand equilibrium, Economic Theory, 5, 19-32.

[5] Dastidar, K., 2006, On the existence of pure strategy Bertrand equilibrium revisited, unpublished manuscript.

[6] Einy, E.: O. Haimanko; D. Moreno; B. Shitovitz, 2009, On the existence of Bayesian Cournot equilibrium, forthcoming in Games and Economic Behavior.

[7] Hoernig, S., 2002, Mixed Bertrand equilibria under decreasing returns to scale: an embarassment of riches, Economics Letters, 74, 359-62.

[8] Kaplan, T. and D. Wettstein, 2000, The possibility of mixed-strategy equilibria with constant-returns-to-scale technology under Bertrand competition, Spanish Economic Review, 2, 65-71.

[9] Mas-Colell, A.; M. Whinston and J. Green, 1995, Microeconomic Theory, Oxford University Press.

[10] Reny, P., 1999, On the existence of pure and mixed strategy Nash equilibria in discontinuous games, Econometrica, 67, 1029-56.

[11] Rudin, W., 1976, Principles of Mathematical Analysis, McGraw-Hill.

[12] Spulber, D., 1995, Bertrand competition when rivals' costs are unknown, Journal of Industrial Economics, 43, 1-11.

[13] Tirole, J., 1988, The Theory of Industrial Organization, MIT Press.

[14] Vives, X., 1999, Oligopoly Pricing: Old Ideas and New Tools, MIT Press.

[15] Weibull, J., 2006, Price competition with convex costs, unpublished manuscript.