Comments Welcome

PQ-Nash Duopoly:

A Computational Characterization

J. Huston McCulloch

March 7, 2011

The author is grateful to participants in the OSU Theory/Experimental brownbag seminar for helpful comments and suggestions.

mcculloch.2@osu.edu

PDF with color figures at <www.econ.ohio-state.edu/jhm/papers/PQNash.pdf>

ABSTRACT

In a duopoly market in which a single good is produced with rising marginal costs and in which each firm must choose both price and quantity produced simultaneously, there is no single-stage Nash equilibrium in pure strategies. In particular, the Cournot, Bertrand, and Stackelberg outcomes are not Nash equilibria of a single-stage game.

However, computational solutions with a finite approximation to the continuous price-quantity (PQ) strategy space find mixed strategy equilibria in which each firm typically has a dominant "regular high price" strategy, accompanied by a string of low-probability "sale-price" strategies. The PQ-Nash duopoly model thus provides a rational model of equilibrium sale prices and price dispersion.

These computational results confirm the unpublished findings of Gertner (1986), that with continuous strategies, a unique mixed equilibrium exists, in which each firm plays a mass point representing the high-price strategy, combined with a continuous distribution over lower prices, with quantity a unique function of price. The computational simulations suggest that this PQ-Nash equilibrium results in prices that are entirely lower than in the Cournot outcome, but entirely higher than in the quasi-competitive Bertrand outcome.

An alternative one-stage game is also considered, in which the good is produced to order after each firm sets both its price and a limit quantity on the number of orders it will accept at its price. Computational results are less well conditioned numerically in this Produce-To-Order case, but some interesting limited results are reported.

1. Introduction

Consider a duopoly market in which two firms produce quantities Q_1 and Q_2 of an identical good with increasing marginal cost schedules $MC_1(Q_1)$ and $MC_2(Q_2)$, and face a decreasing demand function D(P). Consumers are atomistic price-takers with no market power, and price discrimination is not feasible. The two firms offer their output for sale at prices P_1 and P_2. If the two prices are different, the low-price seller sells out first to the highest-demand consumers. Output is perishable, but any unsold output may be disposed of at no cost.

The optimal cooperative strategy for the two firms is either to merge or to collude, and to charge the monopoly price at which combined marginal cost equals marginal revenue. They will then allocate output to the two firms as a "two-plant monopolist" by equating marginal cost in the two "plants." However, it may be that merging or otherwise cooperating is prohibited by antitrust laws, so that they are forced to act as a non-cooperative duopoly. And even if cooperation is permitted, the firms must know what the non-cooperative duopoly outcome would be in order to negotiate a division of the gains from cooperating.

Section 2 below shows that the traditional Cournot, Bertrand, and Stackelberg duopoly models are not Nash equilibria for a single-stage game in which prices and quantities produced are both determined simultaneously. Section 3 shows that there can be no pure Nash equilibria for this game, computes illustrative mixed Nash equilibria, and compares the results to the theoretical findings of Gertner (1986). Section 4 explores an alternative one-stage "production-to-order" game, in which firms set prices and limit quantities simultaneously, and finds mixed equilibria but no pure equilibria unless firms are equal in size. Section 5 briefly discusses Stackelberg models, while section 6 concludes and enumerates unresolved issues.

2. Traditional models of duopoly

In order to provide simple illustrations while still permitting variation in both the relative sizes of the two firms and the elasticity of demand, the present note considers examples in which the marginal cost schedules are linear:

$$\mathrm{MC}_1(Q_1) = Q_1/a,$$

$$\mathrm{MC}_2(Q_2) = Q_2/(1\text{-}a),$$

where $a \in (0, .5]$ is the share of the first, smaller firm in horizontally summed marginal cost. Also in the illustrations, demand is affine in price:

$$\mathrm{D}(P) = 1 + e - eP,$$

where $e > 0$ is the absolute value of the price elasticity at the point $Q = 1$, $P = 1$. The marginal cost and demand schedules are scaled so that $Q = 1$, $P = 1$ is the quasi-competitive equilibrium outcome, at which price equals marginal cost for both firms. In this outcome, $Q_1 = a$ and $Q_2 = (1\text{-}a)$, as shown by the asterisks in Figure 1, which is drawn for $a = 1/3$, $e = 1$. The two-plant monopoly outcome is represented by the + symbols.

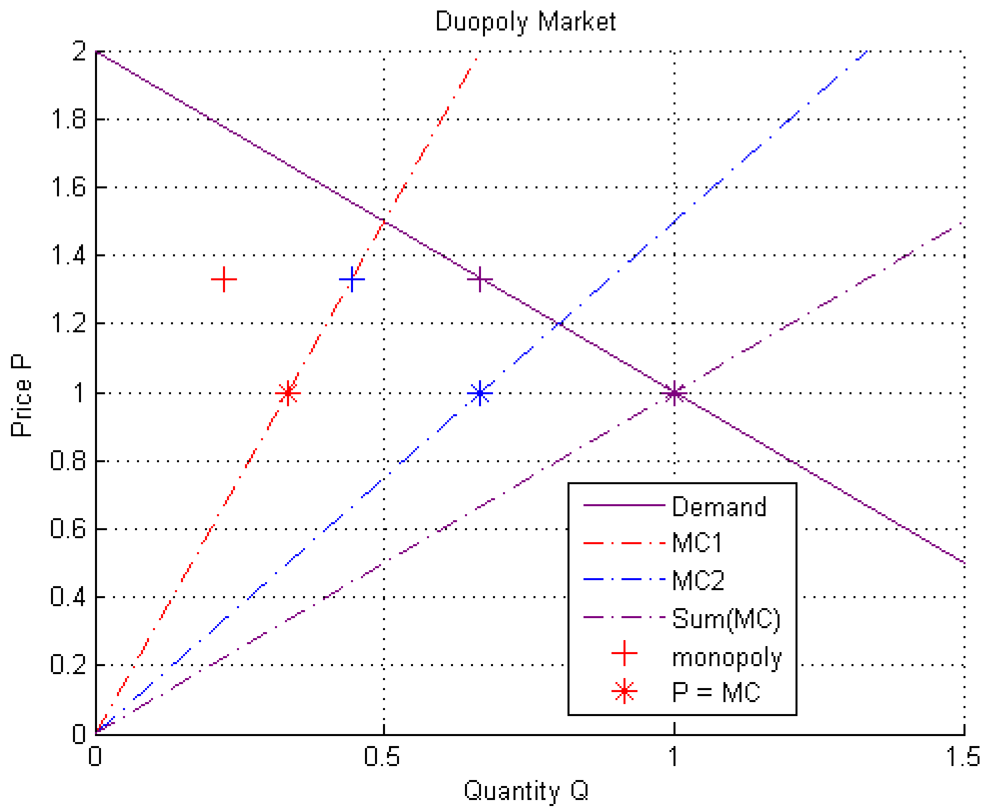

Figure 1

Figure 2 below shows the traditional Cournot, Bertrand, Q-Stackelberg and P-Stackelberg duopoly outcomes in such a market, for the case $e = 1$ and $a = 1/3$.

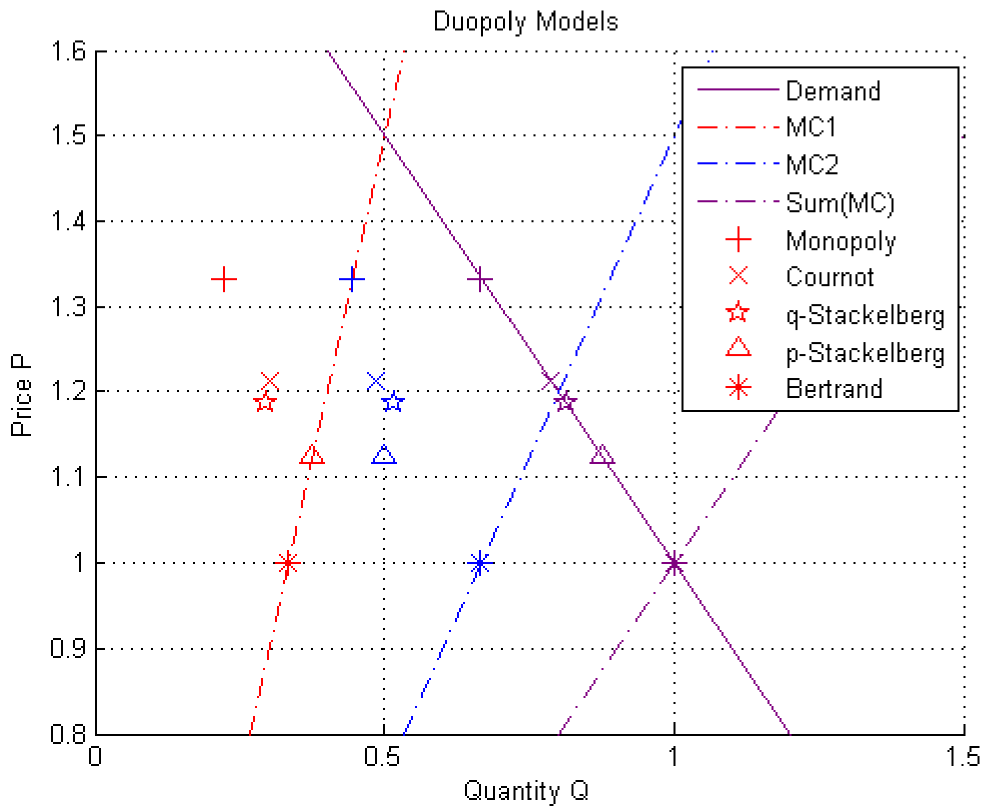

Figure 2

In the traditional exposition of the Cournot (1838) model of duopoly, each firm i assumes the strategy of its rival firm j to be completely determined by its choice of Q_j, and that the rival will sell this output at any price, no matter how low. Firm i then faces a residual demand curve equal to $D_i(P_i) - Q_j$, as shown in Figure 3 below. It selects its own P_i and Q_i so as to maximize its profit with this residual demand curve, by setting marginal cost equal to marginal revenue as computed from this residual demand curve. Firm j does likewise. In equilibrium, the two prices must be equal and quantities must sum to total demand at the common price. If each firm knows the other's marginal cost schedule, both can compute the Cournot equilibrium and move there directly, but even if they did not know the other's marginal cost schedule, they could find this equilibrium by iterative learning over time, so long as each can observe the other's output and they both know the demand curve. The Cournot outcome is represented in the Figures by × symbols.

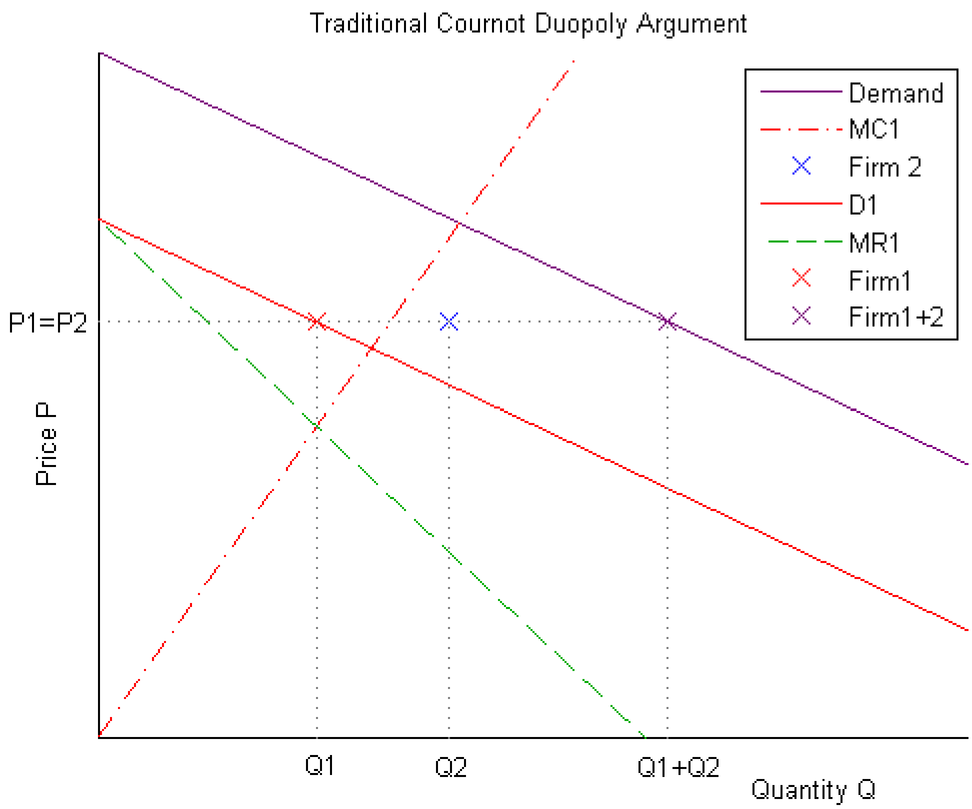

Figure 3

The Cournot model appears to be the equilibrium of what might be called a "Q-Nash" single-stage game, in which each firm's strategy is defined solely in terms of its choice of quantity produced with no regard for price.[1] In fact, however, it is the two

[1] See, for example, Singh and Vives (1984).

duopolists who set price in this economy. If a consumer were to offer either firm \$0.05 for a unit of output, it would refuse, since it will not settle for anything less than the Cournot price (\$1.21 in Figure 2).

Since each firm in fact has a reservation price equal to its offering price as part of its strategy, the Cournot model does not adequately describe the strategy space and therefore is not a true single-stage Nash equilibrium. In fact, firm *i*'s residual demand curve equals $D(P_i) – Q_j$ only down to P_j, at which point it becomes horizontal at price just slightly under P_j, until it reaches the market demand curve, as shown in Figure 4 below, for $i = 1$ and $j = 2$. Because of this kink in firm *i*'s residual demand curve, its marginal revenue schedule has a discontinuity at $Q_i = D(P_j) – Q_j$ at which it jumps from well under P_j up to $P_j - \varepsilon$.[2] Since marginal revenue now exceeds marginal cost for Q_i just above *i*'s Cournot output, firm *i* will charge $P_j - \varepsilon$ and increase its output above the Cournot level. Cournot is therefore not a true Nash equilibrium in a world in which the duopolists in fact set both price and quantity in a single stage game. This analysis suggests that the actual duopoly price will be below the Cournot price.

[2] This kinked duopoly demand curve differs from that proposed by Sweezy (1939), in that it is less elastic above the kink than below. Sweezy supposed that duopolists would be reluctant to match one another's price increases, but would be quick to match price cuts.

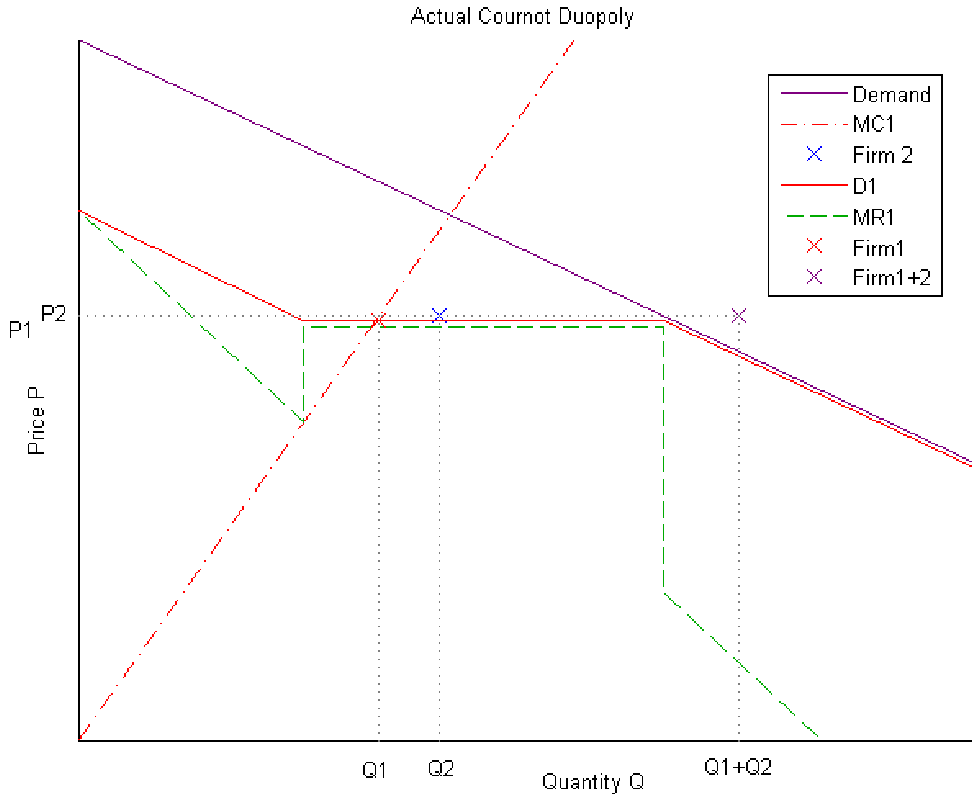

Figure 4

Kreps and Scheinkman (1983) demonstrate that Cournot *is* the equilibrium of a *two-stage* game in which both firms set their quantity produced in the first stage, and then, knowing each other's production, set their prices in the second stage. However, if the firms choose quantity produced and price simultaneously, or even if they produce first but don't know the other's production until after prices have been chosen, we are back to a *one-stage* game in which Cournot is not a true Nash equilibrium. The present paper is concerned with the more difficult single-stage duopoly game.

In the rival Bertrand (1883) model of duopoly, each firm is instead assumed to take its rival's *price* as its strategy, and to assume that the rival is willing to sell any quantity at this price. It then sets its own quantity as $Q_i = \mathrm{MC}_i^{-1}(P_j)$, slightly undercutting P_j if necessary in order to sell this output. However, at any price above the quasi-competitive price $P = 1$, one or both firms will need to cut price until $P = 1$ is reached. The Bertrand equilibrium is therefore the quasi-competitive equilibrium in which price equals marginal cost for both producers, as shown in Figure 5. The Bertrand outcome is represented in the Figures by the asterisk symbols.

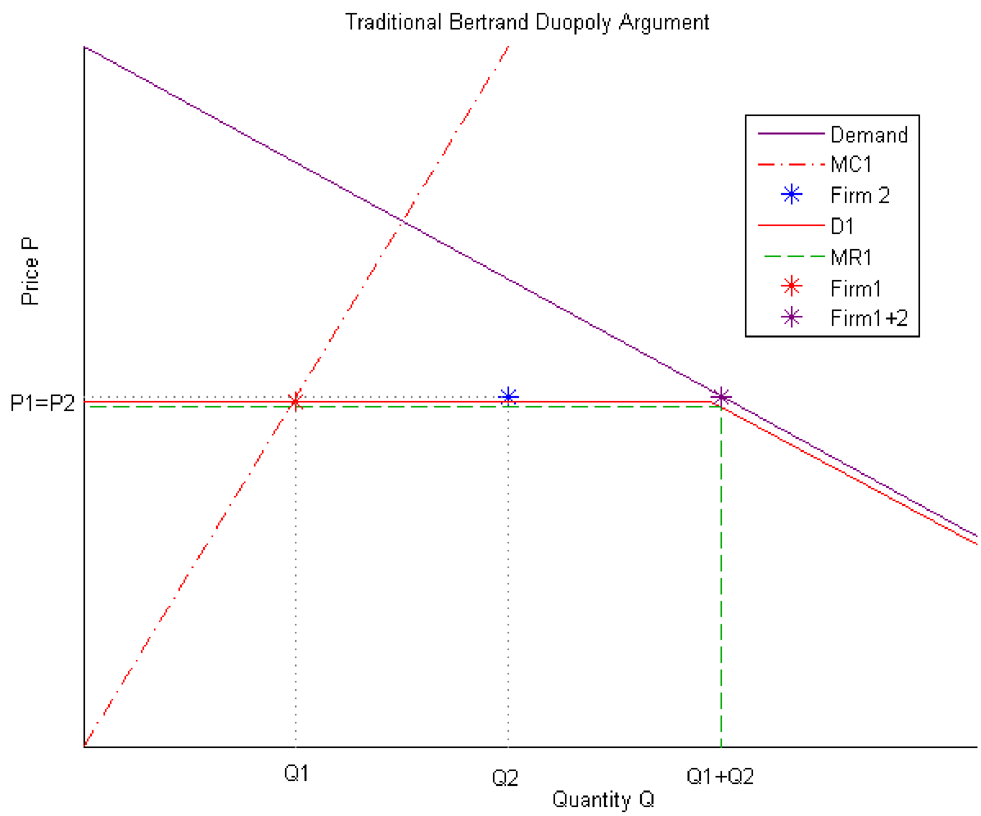

Figure 5

The Bertrand model appears to be what might be called a “P-Nash” equilibrium, in which each firm’s strategy is defined entirely in terms of its price.[3] In fact, however,

[3] See, e.g., Singh and Vives (1984).

each firm carefully selects both quantity and price, and would not meet unlimited demand at the equilibrium price, as falsely assumed by its rival, so long as its marginal cost rises without limit. Once again each firm has a kinked demand curve, as shown in Figure 6 below for firm 1, but now since marginal revenue is below marginal cost for quantities below the Bertrand outputs at the Bertrand price, firm 1 tries to cut its output, forcing P_1 up above the Bertrand price. At this price, the market clears, but firm 2 now wishes it had charged a higher price as well. Therefore Bertrand is not a true Nash equilibrium either. This analysis suggests that the true duopoly price will be above the Bertrand price.

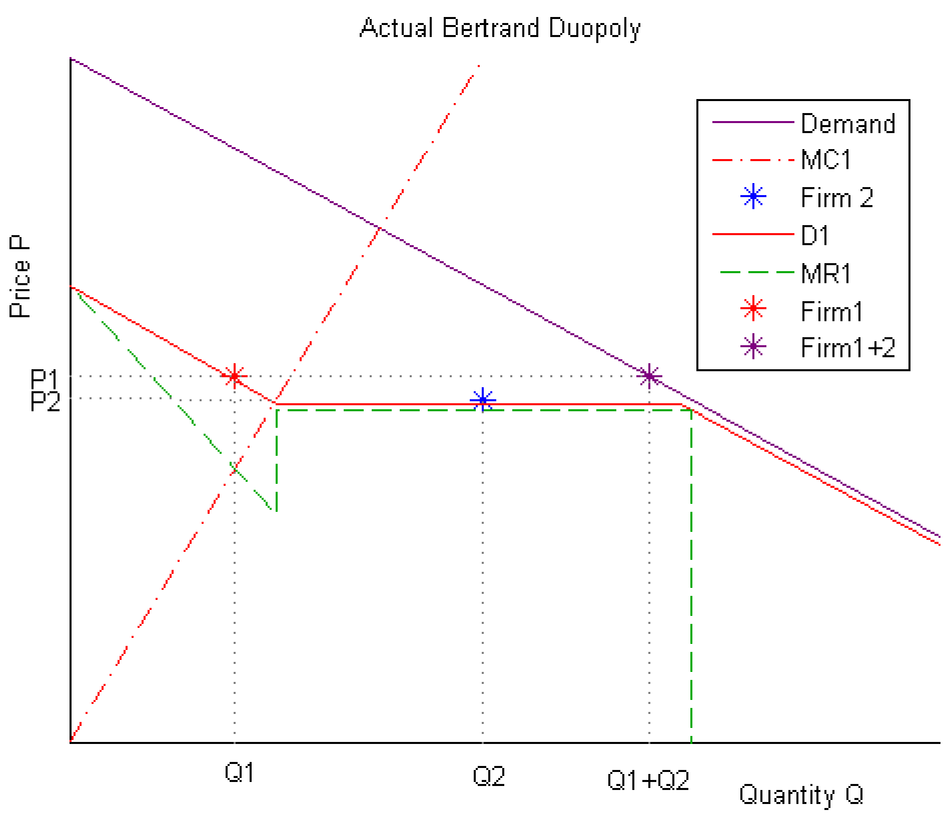

Figure 6

3. PQ-Nash Duopoly

The obvious solution is to model *both* firms as having a strategy defined in terms of *both* price and quantity produced, in what might be called a *PQ-Nash Duopoly* equilibrium.

A problem immediately arises, however, in that there can be no Nash equilibrium in terms of pure strategies in such a market: As discussed above, if firm j sets P_j and Q_j, firm i will face a decreasing residual demand curve $\mathrm{D}(P_i) - Q_j$ down to $P_i = P_j$, and then a flat demand curve at $P_j - \varepsilon$ out to $\mathrm{D}(P_j)$. Its marginal revenue schedule will lie under the decreasing leg of the demand curve, and then discontinuously rise to $P_j - \varepsilon$ at the kink in the demand curve. If its marginal cost schedule just intersects this marginal revenue schedule on its decreasing leg at the bottom of the discontinuity, as at the Cournot price and quantities in Figure 4, there will be a second, higher profit intersection at some $Q_i >$ $\mathrm{D}(P_j) - Q_j$ that it will prefer to move to, so that quantity supplied will exceed quantity demanded and this could not be an equilibrium. The situation is even worse if firm i's marginal cost schedule only intersects marginal revenue once, to the right of the discontinuity.

Or, if i's marginal cost schedule just intersects marginal revenue at price P_j and quantity $\mathrm{D}(P_j) - Q_j$, as at the Bertrand price and quantities in Figure 6, there must be a second, higher profit intersection to the left of this point that it will prefer to move to. But at such a point, quantity supplied will fall short of quantity demanded, and so this cannot be an equilibrium either. Again, the situation is even worse if marginal cost does not intersect the horizontal portion of the marginal revenue schedule at all, but lies entirely to the left of it.

The only other possibility (assuming, for simplicity, that the decreasing leg of the residual demand curve implies a monotonic decreasing marginal revenue schedule as in our linear example) is that there are *three* intersections with marginal revenue, one to the left of the discontinuity, one at the discontinuity itself, and one to the right of the discontinuity on the horizontal portion of the residual demand curve, as shown in Figure 7 below.[4] In this case, the central intersection clears the market, but unfortunately, this is a profit *minimum* rather than a profit *maximum*. The two other intersections are local profit maxima, but neither of them clears the market. If both of them happen to give equal profits, firm 1 will be indifferent between them, and would also be willing to select between them randomly with any probabilities p and $1-p$, so that there would not be unambiguously excess demand or supply. At this point, however, we are out of the realm of pure strategies and into the realm of mixed strategies. *There is therefore no single-stage Nash equilibrium in pure strategies when both firms choose their prices and quantities produced simultaneously.*

[4] Figure 7 is in fact drawn with firm 2's price and quantity equaling the average of its Cournot and Bertrand prices and quantities.

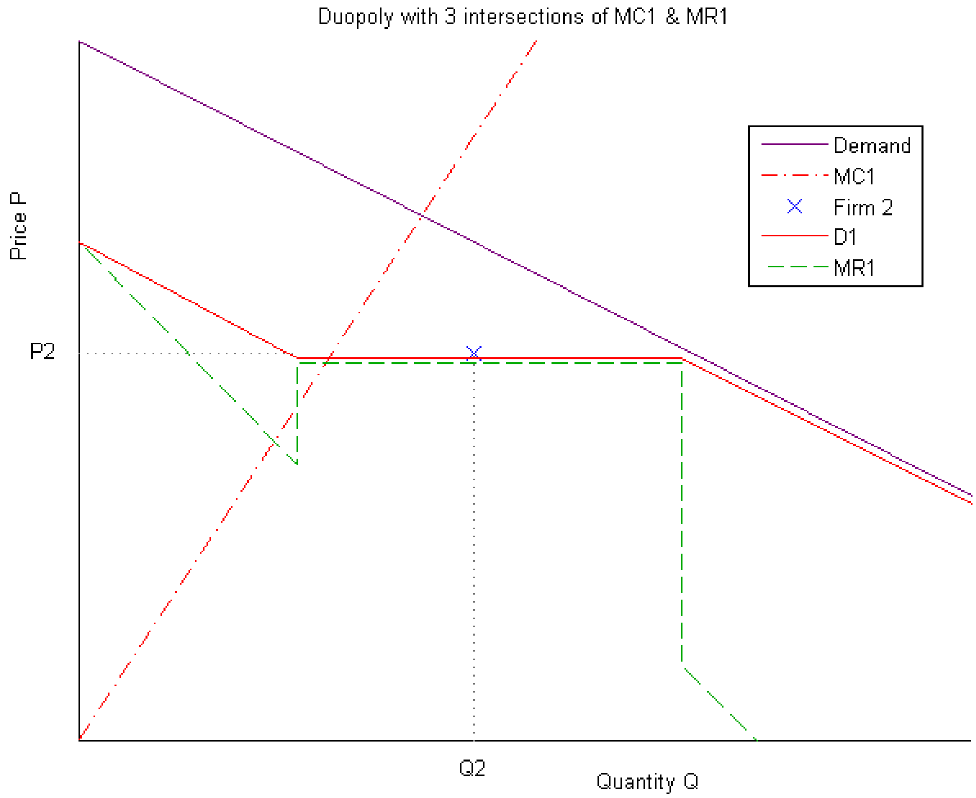

Figure 7

Although there is no PQ-Nash duopoly equilibrium in pure strategies, the famous theorem of Nash (1950) fortunately states that in a finite bimatrix game there must always be at least one *mixed* strategy Nash equilibrium. Although with continuously variable prices and quantities the strategy space of each firm is infinite, it can be arbitrarily well approximated by a discrete approximation that is guaranteed at least one mixed Nash equilibrium.

With mixed strategies, there is some chance that combined output will exceed or fall short of demand at the higher price. Let q_i be the amount actually sold by firm i, as contrasted with Q_i, which continues to represent the amount produced. Then if $P_i < P_j$, it is assumed that the low priced firm sells out first to the highest demand customers:

$$q_i = \min(Q_i, \mathrm{D}(P_i)),$$

$$q_j = \min(Q_j, \max(0, \mathrm{D}(P_j) - q_i)).$$

In the rare event that both firms charge exactly the same price $P = P_i = P_j$ and there is excess supply, we adopt the reasonable "tie-breaker" rule that sales are allocated in proportion to production:

$$q_i = \min(Q_i, \mathrm{D}(P)\, Q_i / (Q_i + Q_j)).$$

Any unsold output is assumed to be perishable and hence discarded, with zero disposal cost. The profit of firm i is then

$$\Pi_i = P_i\, q_i - \mathrm{C}_i(Q_i) = P_i\, q_i - Q_i^2 / (2a_i),$$

where $a_1 = a$ and $a_2 = 1\text{-}a$.

It is not clear *a priori* what the nature of the mixed Nash equilibrium should be: Will each firm choose between just two strategies, a high "regular" price with low output and a low "sale" price with high output? In this case, each mixed strategy would be a probability mass function with only two mass points. The two prices could either be the same for both firms, in which case the tie-breaker rule will be relevant, or, as is more likely, one firm (presumably the smaller) will slightly undercut the other by setting its prices just slightly lower than the other's, in order to guarantee itself a sell-out when both firms choose a similar price.

Or will each firm's mixed strategy have a continuous bivariate distribution whose support is a subset of the region under the demand curve? In this case, the probability of a tie would be zero and hence have no effect on expected profits.

Or will each firm's strategy have a continuous distribution on a one-dimensional manifold of prices and quantities? If so, is the manifold described by a monotonic function of either price or quantity, or can it bend back on itself in either dimension?

Or are the strategies mixtures of one or more mass points and a continuous distribution?

In order to shed some light on these questions in a variety of cases, Figures 8 – 12 show computational approximations to mixed PQ-Nash equilibria for e = 1, 2 and 1/2 with a = 1/3, and also for e = 1 with a = 1/10 and 1/2. A fairly fine 15×15 rectangular grid of price and quantity combinations for each firm was used to approximate the continuous strategy space, giving each firm $15^2 = 225$ potential strategies. On the assumption that the continuous-space equilibrium is a 1- or 2-dimensional continuous distribution with zero probability of ties, the two firms' potential price values were offset by 1/2 step in order to preclude ties altogether. A $15^2 \times 15^2 = 225 \times 225$ payoff matrix was then generated for each firm, and one mixed Nash equilibrium was computed using an algorithm based on Mangasarian and Stone (1964). The areas of the red squares are proportional to the probability of each (P, Q) strategy combination for firm 1, and the areas of the blue diamonds are proportional to these probabilities for firm 2. The squares and diamonds in the legend boxes are scaled to represent probability 1/4. For comparison, the Cournot and Bertrand equilibria are also plotted. Further computational details are in the Appendix.

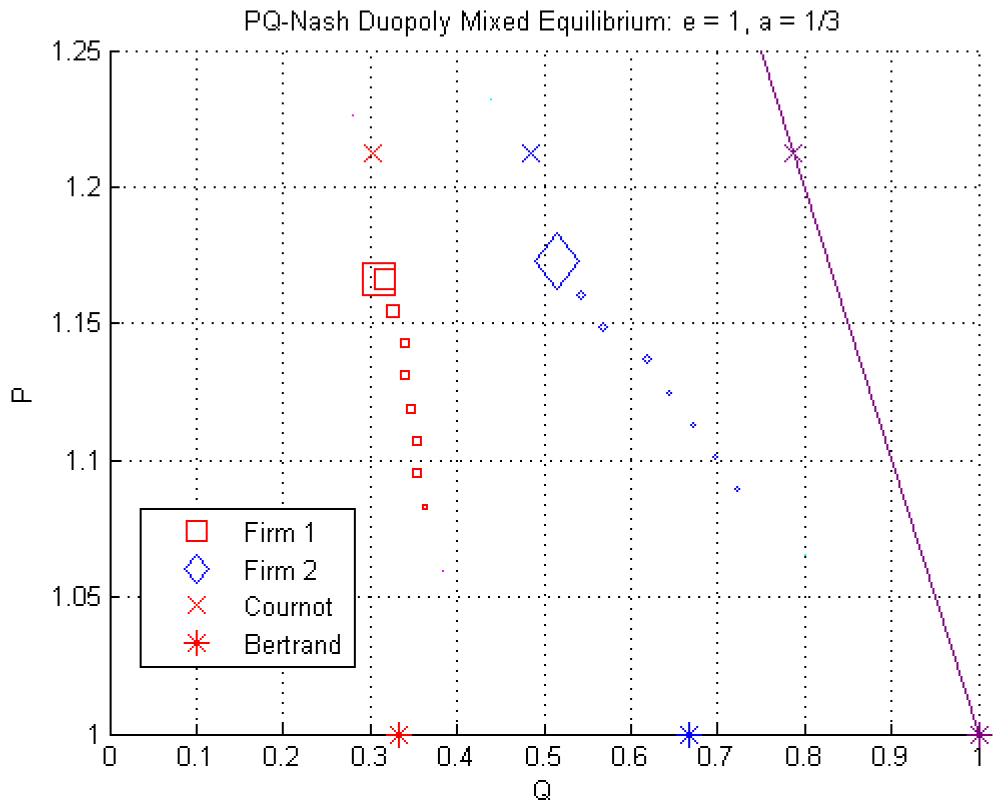
PQ-Nash Duopoly Mixed Equilibrium: e = 1, a = 1/3
P
Q
1.25
1.2
1.15
1.1
1.05
1
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
Firm 1
Firm 2
Cournot
Bertrand

Figure 8

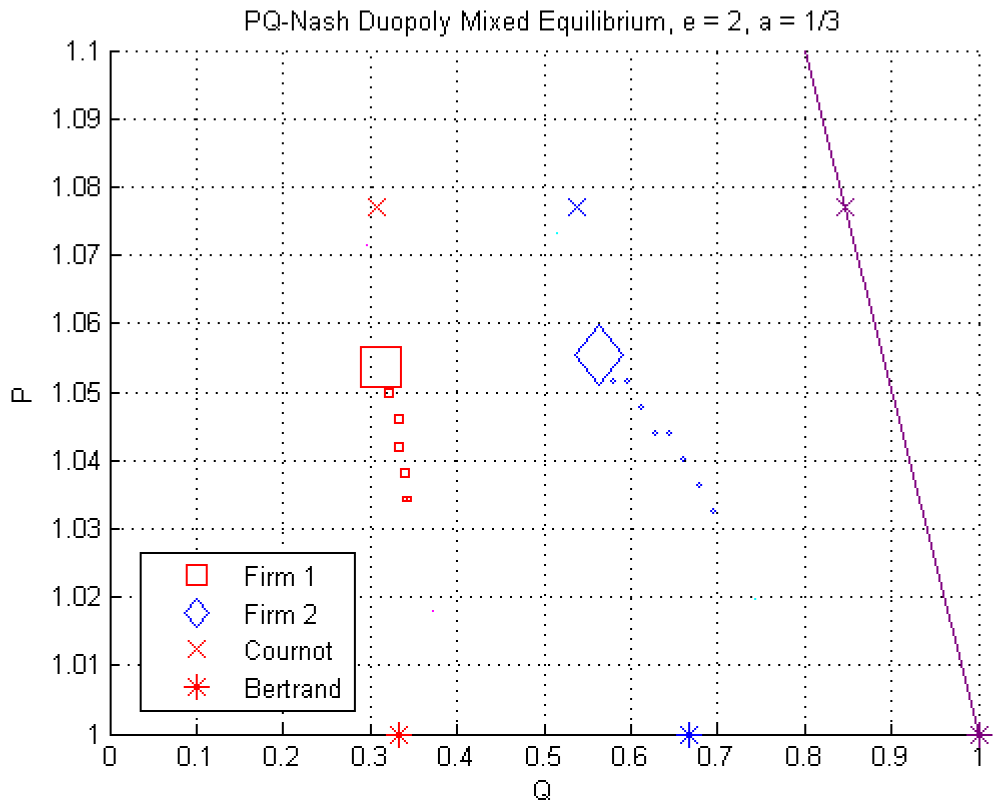
PQ-Nash Duopoly Mixed Equilibrium, e = 2, a = 1/3
P
Q
Firm 1
Firm 2
Cournot
Bertrand
1.1
1.09
1.08
1.07
1.06
1.05
1.04
1.03
1.02
1.01
1
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

Figure 9

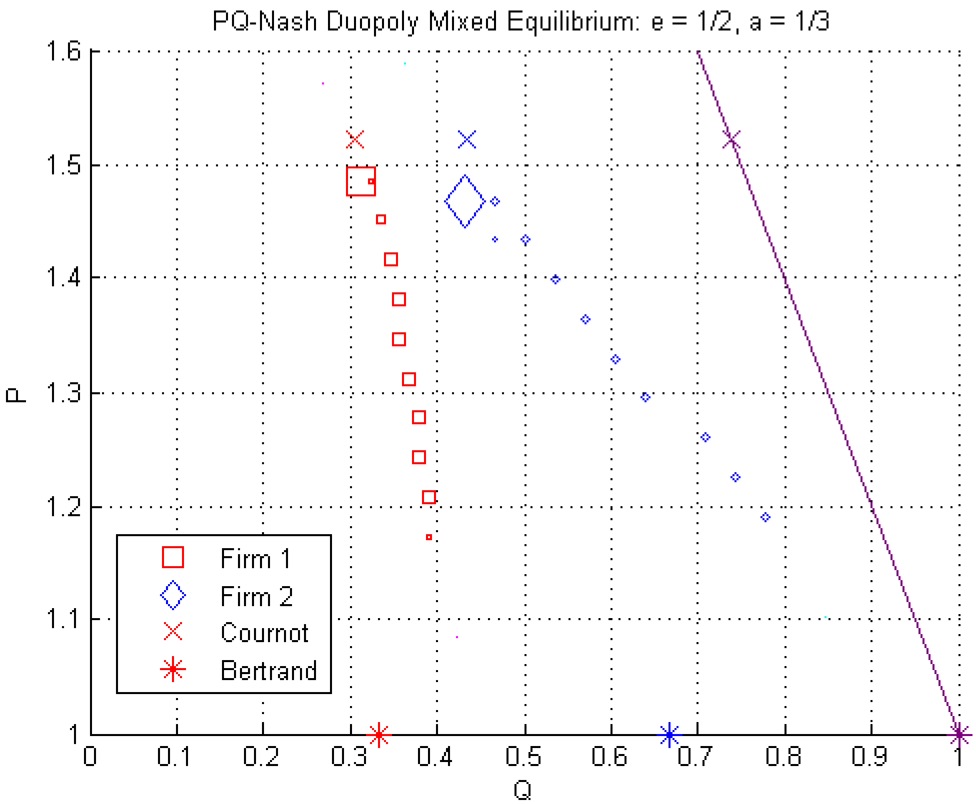
PQ-Nash Duopoly Mixed Equilibrium: e = 1/2, a = 1/3
P
Q
1.6
1.5
1.4
1.3
1.2
1.1
1
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
Firm 1
Firm 2
Cournot
Bertrand

Figure 10

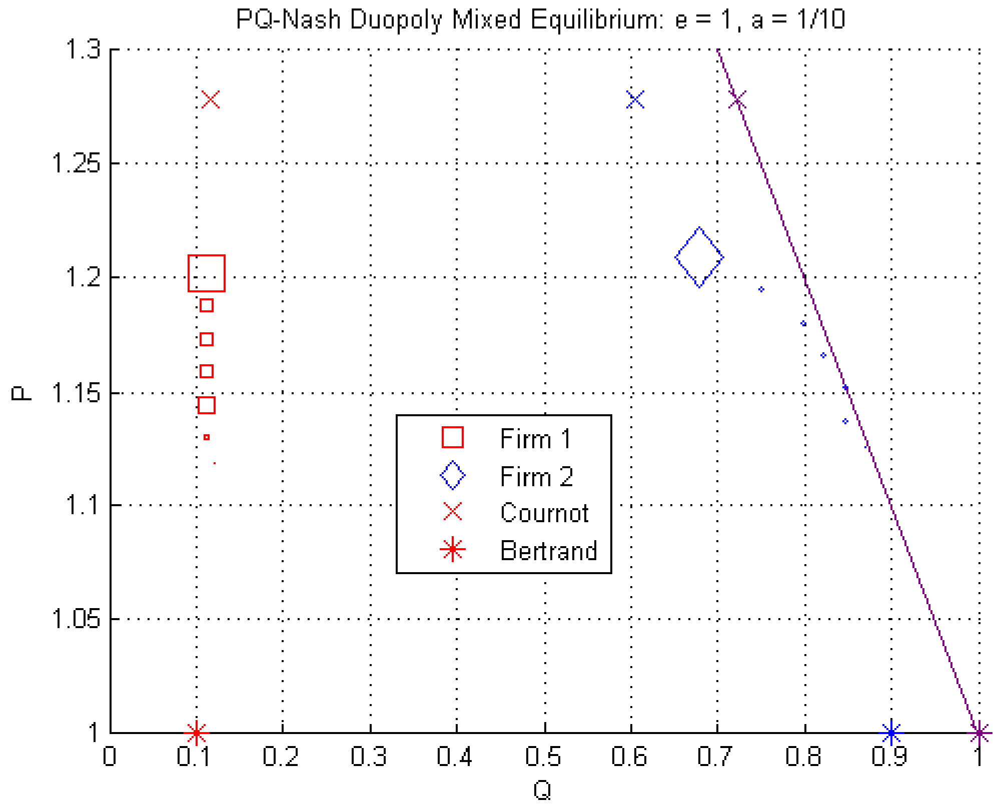
PQ-Nash Duopoly Mixed Equilibrium: e = 1, a = 1/10
1.3
1.25
1.2
1.15
1.1
1.05
1
P
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
Q
Firm 1
Firm 2
Cournot
Bertrand

Figure 11

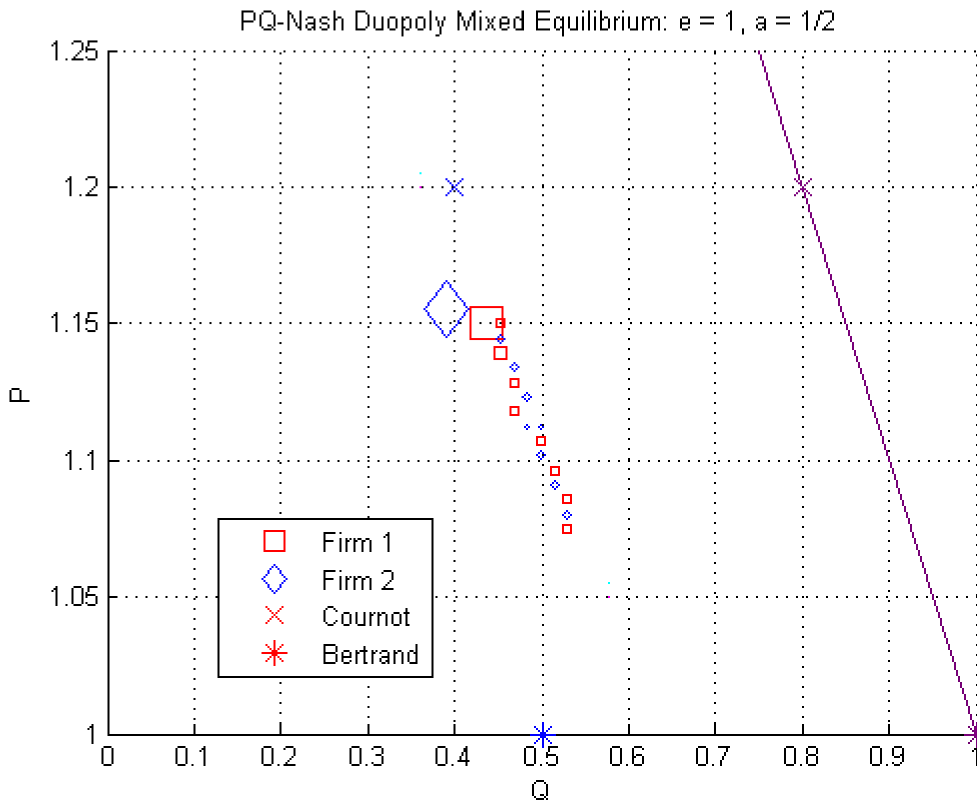

Figure 12

In every case, each firm has a dominant strategy with probability in the range 0.44 to 0.93 that involves a "regular high price" and low quantity, accompanied by a string of small probability strategies with lower "sale prices." For the smaller firm 1, the associated sale quantities are nearly constant, while for the larger firm, they increase briskly as price falls. Each pair of payoff matrices was searched for pure Nash equilibria, and as expected, none were present.

In each case, the PQ-Nash equilibrium is "more competitive," i.e. has unambiguously lower price with generally higher output, than Cournot. Also in each case, the PQ-Nash equilibrium is “less competitive,” i.e. has unambiguously higher price with generally lower output, than Bertrand.

Since the above simulations were performed, I have learned that Robert H. Gernter (1986) has solved the first order conditions for the PQ-Nash duopoly problem, in which price and quantity are determined simultaneously. He considers the case of increasing marginal cost, as well as constant and decreasing marginal cost, for the case of two independent firms with identical marginal cost schedules.

With increasing marginal costs, Gertner finds that each firm has a distribution over prices bounded strictly below by the competitive price and strictly above by the monopoly price. Quantity is a nonrandom function of price, so that the joint distribution is on a one-dimensional manifold. Each firm plays its highest price with positive probability mass, and the market clears when both firms play this strategy. The remainder of the distribution is continuous, and generally leads to unsold output. Gertner’s characterization thus generally confirms the computational results above.

Gertner finds an explicit solution for the case of constant marginal cost, but unfortunately is unable to find an explicit solution with increasing marginal cost, even in the case of straight line demand and marginal cost curves. Also, he considers only the symmetrical case in which the two firms share the same marginal cost schedule.

Unfortunately his essay, which was the second of three essays in his MIT dissertation, has never been published. Nevertheless, a copy may be obtained through MIT’s D-Space service.[5]

---

[5] Tirole (1990: p. 233) does mention this essay, but without providing any details or discussion of the crucial case of increasing marginal cost. Tirole was one of Gertner’s dissertation advisors.

In a repeated pure Nash equilibrium, the players often can learn the equilibrium through a process of iterative learning, even if they do not know the other's payoff matrix: One firm adopts a strategy, the other reacts optimally, and then the first revises optimally, etc. until a steady state is reached.

However, there is no such process by which the players could learn a mixed equilibrium, even with repeated play, without knowing the other's payoff matrix and actually computing the equilibrium. This occurs because in a mixed equilibrium, each player is indifferent over all the strategies to which it assigns non-zero probabilities, and therefore does not care what its *own* probabilities are. All that determines player *i*'s probabilities is the mathematical restriction that they must make player *j* just indifferent over *its* non-zero strategies.

In the duopoly context, however, both firms are in the same industry, and therefore could reasonably know the other's cost function and therefore payoff function. In the examples provided in this paper, in fact, their cost schedules are identical up to a scale factor.

Levitan and Shubik (1978) anticipate Gertner by finding an analytical solution to a mixed strategy PQ-Nash duopoly equilibrium, but only for an economy in which marginal cost is zero or constant, but disposal (or storage to the next period) is costly.

Likewise, Kreps and Scheinkman (1983) anticipate the Gerter analysis in the special case of a vertical marginal cost schedule.[6]

4. A Produce-to-Order Duopoly

A further economically relevant duopoly structure is one in which each firm i sets its price P_i, and then *passively takes orders* up to some *limit quantity* $Q_i^{\max}$ that it sets at the same time as its price. Afterwards, it faithfully produces the quantity ordered, but since no second-stage decision is involved, this is still a single-stage game. As in Section 3 above, the firm makes a simultaneous price and quantity decision, but in this case the quantitative choice variable is $Q_i^{\max}$, rather than *actual* production Q_i.

In this Production-to-Order (PTO) duopoly, actual production Q_i always equals sales q_i, so that there never is overproduction, though there could occasionally be underproduction in a mixed equilibrium. If $P_i < P_j$, we have

$$q_i = \min(Q_i^{\max}, \mathrm{D}(P_i)),$$

$$q_j = \min(Q_j^{\max}, \max(\mathrm{D}(P_j) - q_i, 0)).$$

In the rare event that $P_i = P_j = P$ and the sum of the limit quantities exceed total demand, we need a tiebreaker assumption. However since the limit quantities, unlike quantities already produced, are not observable by consumers until they become binding, there is no reason for consumers to purchase in proportion to them. Since the firms appear identical to consumers, we therefore assume that unless one or both limit quantity is binding, each firm will get an equal share of total demand:

---

[6] Singh and Vives (1984) consider an economy in which firms are artificially constrained to choose between setting their prices and then being required to supply whatever quantity is demanded at that price, or setting their quantities and then having to sell at whatever price clears the market. The present paper instead allows them to set both price and quantity.

$$q_i = \min(Q_i^{\max}, \max(D(P)/2, D(P) - Q_j^{\max})).$$

In our linear example, profit is then

$$\Pi_i = P_i\, q_i - C_i(q_i) = P_i\, q_i - q_i^2 / (2a_i).$$

By setting $Q_j^{\max}$ sufficiently high, firm *j* can push the kink in firm *i*'s residual demand curve so far to the left that it becomes irrelevant, and thereby induce firm *i* to act like a P-Stackelberg follower.

Figure 13 below provides a computational discrete approximation to this PTO duopoly for our benchmark case e = 1, a = 1/3. No offset was incorporated into the two price grids, so that ties occasionally do arise. Note that the horizontal axis now represents the strategic limit quantities, rather than quantities actually produced and sold. No pure strategy Nash equilibria were present with these parameters.

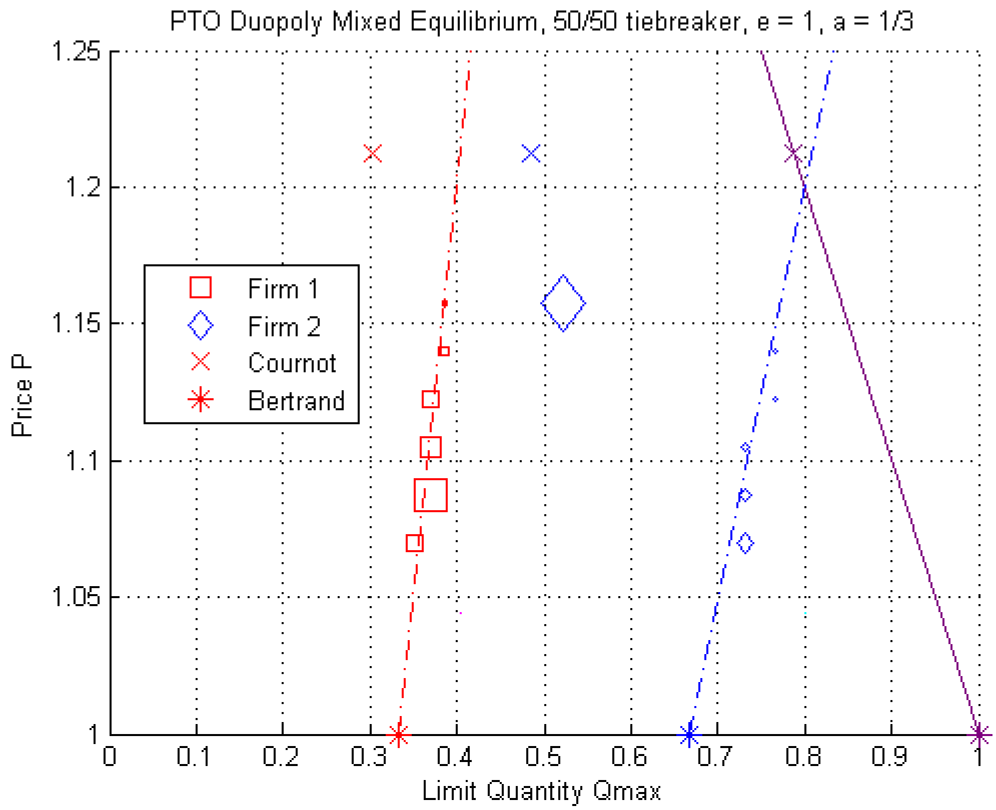

Figure 13

The primary strategy (with probability 0.72) of the larger firm 2 has the highest price charged by either firm, and appears to be a mass point. Firm 2 also has what appears to be a continuum of lower sale prices, but now with the highest density at or near its lowest price, rather than at the high end as in the PQ duopoly. Firm 2's sale prices are all on (or very close to) its marginal cost schedule, while its regular high price is well to the left of its MC schedule.

The primary strategy of the smaller firm 1 (with probability 0.49 in this discrete approximation) is now near the bottom of its price distribution, and may merely represent a mode in a continuous density rather than a mass point. There is no evidence of a mass point at the highest price in the support. Since firm 1 is almost always the low-price vendor, all of its quantities lie on (or very close to) its marginal cost schedule.

Note again that in a mixed strategy, each player is indifferent over all strategies to which it assigns positive probabilities or densities.

When both firms are the same size (a = 1/2) with production-to-order and e = 1, multiple pure-strategy Nash equilibria do appear when both firms are symmetrically given the same price grid with no offset. These vary in price from $P_1 = P_2 = 1.000$ with multiple non-binding limit quantities, up to $P_1 = P_2 = 1.040$ with multiple non-binding limit quantities, so that each firm gets half of demand. The high price equilibria turn out to be what might be called "superlative pure Nash Equilibria", in that any one of them is the optimal Nash equilibrium in pure strategies for both firms, and therefore the one that they could be expected to select when each knows the other's payoff matrix.

The production-to-order duopoly often encountered computational difficulties with other values of the parameters, perhaps due to the relatively indeterminate value of the limit quantity in some cases. Accordingly, it has not been explored as thoroughly as the PQ case, discussed in Section 3, in which quantity actually produced is the strategic variable.

However, it is conjectured that when one firm is smaller than the other and the strategy space is continuous, there will be no pure Nash equilibrium, yet there will be a unique mixed equilibrium in which the larger firm primarily charges a high price with restricted limit quantity, plus a string of lower prices, with limit quantities on or near its marginal cost schedule, while the smaller firm charges a continuum of low prices, with limit quantities that are at or near its marginal cost schedule. When both firms are equal in size, it is conjectured that there will be a superlative Nash equilibrium in pure strategies in which both firms charge a moderately high price, with relatively high limit quantities that are non-binding.

5. Stackelberg outcomes

In Stackelberg models, one firm, presumably the larger firm 2 in our example, is assumed to be the “market leader,” and the other, presumably smaller firm, to be the “follower.” Two variants of the Stackelberg model are illustrated in Fig. 2. In what might be called the “Q-Stackelberg” variant, the follower takes the leader’s *quantity* as given, and then chooses price and quantity so as to maximize its own profits under the assumption that the leader will sell this quantity at any price as in Cournot. The leader then subtracts the follower’s supply from the total demand curve at each price, and maximizes its own profits by choosing a price and quantity on this residual demand curve. This gives a well-defined outcome, indicated in Figure 2 by the pentagram symbols. However, this is not a Nash equilibrium, since the moves are not simultaneous. Furthermore, the follower firm wrongly assumes the leader’s strategy consists only of a quantity decision, when in fact the leader is choosing both price and quantity.

Nevertheless, Q-Stackelberg can be justified as a 3-step game, in which the leader first chooses its quantity, then the follower chooses its quantity, and then prices are determined in a third, Nash step as in Kreps and Scheinkman (1983).

In what might be called the “P-Stackelberg” variant, the follower instead takes the leader’s *price* as given and then chooses price and quantity so as to maximize its own profits under the assumption that the leader will sell any amount at this price. The leader then subtracts the follower's supply (which now just equals $MC_1^{-1}(P_2)$) from total demand and again maximizes its own profits by choosing a price and quantity on the residual demand curve, as indicated by the triangle symbols in Figure 2. Again this is not a true Nash equilibrium, however, since the follower firm now wrongly assumes the leader's strategy consists only of a *price* decision.

P-Stackelberg is perhaps a good approximation to duopoly equilibrium when a is near zero, so that the follower firm 1 is in fact facing a nearly infinitely elastic residual demand curve to the left of $Q = D(P_2)-Q_2$. However, as a increases above zero, this elasticity becomes perceptibly finite, so that firm 1 wants to cut its quantity and raise its price above these values. Even if firm 1 persists in believing that firm 2 would sell any quantity at P_2, a further problem arises as a becomes larger, in that eventually, even short of $a = 1/2$, neither firm wants to the leader, since the leader must bear the entire burden of cutting output so as to hold price up for the benefit of both firms. If it is arbitrarily assumed that the larger firm is the leader, the outcome will change discontinuously as a passes 1/2, a problem that arises in McCulloch (1993).

The inconsistency in a 2-step Q- or P-Stackelberg game could be resolved by assuming the follower takes both the leader's price *and* quantity into account, in what might be called a *PQ-Stackelberg* game (not illustrated). Since by adjusting its quantity, the leader can induce the follower to be either a price-taker as in Figure 4 or a quantity-taker as in Figure 6, the follower is always acting as either a Q-Stackelberg or P-Stackelberg follower. At the boundary between the two zones, there are three intersections of MC and MR as in Figure 7 with equal profit in the left and right intersection, so that the boundary is easy to locate. It is therefore not inconceivable that the outcome will always be either Q- or P-Stackelberg, and perhaps even always one or the other, unless induced zone changes affect the leader's optimal strategy.

In any Stackelberg model, a big unsolved issue is which firm goes first. However, the decision to go first or wait is in fact itself a Nash game: Perhaps the NE is that one firm (the larger?) decides to go first while the other agrees to wait, or perhaps it is that both try to go first and therefore move together. Or perhaps this game will itself require a mixed strategy solution.

6. Conclusions and unsolved issues

The Gertner PQ-Nash model solves the 127-year-old debate over whether Cournot or Bertrand had the better model of duopoly: Although Cournot can be justified as a two-stage process in which first quantities are simultaneously chosen and revealed and then prices are simultaneously chosen, *neither is correct* as a single-stage determination of both prices and quantities. As Bertrand rightly noted, Cournot does not adequately take price competition into account. However, Bertrand was also wrong, not to take quantity restriction into account.

When both price competition and quantity restriction are modeled as a one-stage game, the Gertner mixed strategy game arises with a price support intermediate between Cournot and Bertrand. The resulting mixed strategy game provides a rational model of sales and price dispersion.

It remains to find an explicit or at least easily computable solution for Gertner's differential equations for simple parametric cases. In particular, linear marginal cost and affine demand should make quantities an affine function of price, and yield a closed form solution. Constant elasticity demand with Cobb-Douglas or CES marginal variable cost schedules (arising from diminishing marginal product against a fixed second factor) are of even greater interest. These probably do not lead to an explicit solution, but there should be a quick and accurate algorithm to approximate the solution arbitrarily closely.

Gertner (1986) only treats the symmetric case of equal marginal cost schedules. It should be straightforward to generalize this to the case of different marginal cost schedules and then to solve the resulting differential equations for the linear/affine case, and to find an algorithm to approximate more general cases.

The Production-to-Order dupoly of Section 4 and the PQ-Stackelberg duopoly model of Section 5 also remain to be solved. It could be that the PQ-Stackelberg model provides a good approximation to the PQ-Nash model, at least in the case of unequal firm sizes, yet may be easier to compute.

And finally, it should be straightforward (so to speak) to extend the Gertner and other models of duopoly discussed here to the case of an n-firm oligopoly. Presumably the maximum price will converge on the competitive price at approximately the same rate as in the original Cournot model.

REFERENCES

Bertrand, J. (1883). "Théorie des Richesses" (Review of *Théorie Mathématique de la Richesse Sociale* by Léon Walras and *Recherches sur les Principes Mathématiques de la Théorie des Richesses* by Augustin Cournot), *Journal des Savants* (Sept.), pp. 499-508.

Chatterjee, Bapi (2008). Matlab script *bimat.m*, online at Matlab Central File Exchange, <www.mathworks.com/matlabcentral/fileexchange/22356-bimatrix-game>.

Cournot, Augustin (1838). *Recherches sur les Principes Mathématiques de la Théorie des Richesses*, Paris. (Not consulted)

Gertner, Robert H. Simultaneous move price-quantity games and equilibrium without market clearing. Essay Two of R.H. Gertner, *Essays in Theoretical Industrial Organization*, Doctoral Dissertation, MIT, 1986. Available through MIT's D-Space service, <http://dspace.mit.edu/handle/1721.1/14892>

Kreps, David M., and José A. Scheinkman (1983). Quantity precommitment and Bertrand competition yield Cournot outcomes. *Bell Journal of Economics* **14**: 326-337.

Levitan, R., and M. Shubik (1978). Duopoly with price and quantity as strategic variables. *International Journal of Game Theory* **7**: 1-11.

Mangasarian, O.L. (1964). Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics* **12**: 778-780.

Mangasarian, O.L., and H. Stone (1964). Two-person nonzero-sum games and quadratic programming. *Journal of Mathematical Analysis and Applications* **9**: 384-355.

McCulloch, J. Huston (1993). Is antitrust beneficial? Increasing returns to scale with sunk costs." Ohio State University Working Paper, online at <www.econ.ohio-state.edu/jhm/papers/Antitrust.pdf>.

Nash, John F., Jr. (1950). Equilibrium points in N-person games. *Proceedings of the National Academy of Sciences* **36:** 48-49.

Singh, Nirvikar, and Xavier Vives (1984). Price and quantity competition in a differentiated duopoly. *Rand Journal of Economics* **15**: 546-554.

Sweezy, Paul M. (1939). Demand curves under conditions of oligopoly. *Journal of Political Economy* **4:** 58-73.

Tirole, Jean (1988). *The Theory of Industrial Organization.* Cambridge: MIT Press. (1990 printing consulted.)

Appendix:

Computational Considerations

Numerical mixed equilibria were found using a modification of the very helpful function *bimat.m*, written in Matlab by Bapi Chatterjee (2008), and contributed by him to Matlab Central File Exchange. This program is based on the Mangasarian and Stone (1964) quadratic programming algorithm, and relies on the intrinsic Matlab function *quadprog.m*.

There is, however, an important bug in Chatterjee's function, in that he neglected to check solutions for the Mangasarian and Stone requirement that the value of the quadratic program be zero (or at least be within computational error of zero). Accordingly, a similar function *bimatrix.m* was written by the author, which tries different initializations until a solution is found that meets this condition. Two other minor bugs were also corrected.

Each player is given an $h \times h$ rectangular grid of prices and quantities, giving it h^2 strategies, and making the payoff matrices $h^2 \times h^2$. The matrix passed to *quadprog.m* is then $2h^2 \times 2h^2$ Since with an $n \times n$ matrix, *quadprog.m* typically requires about 2 or 3 times n iterations to find even an approximate solution, and since each iteration requires on the order of n^2 floating point operations, the theoretical total computation time is roughly proportional to h^6. In practice, it was found that $h = 11$ (with 10 equal steps) takes several seconds on a laptop, $h = 15$ takes several minutes, and $h = 21$ takes several hours. Accordingly h was restricted to 15. The maximum iterations for *quadprog.m* were conservatively set to $10h^2$ or 2250, but this was never binding.

In order to provide a fairly fine grid in the relevant region without excessive computational time, a two-step procedure was adopted: In the first step, if no price offset was used, each firm's minimum price was set to the Bertrand price, and each firm's maximum price was set to the Cournot price. When a price offset was used, the range was widened slightly so that neither firm's minimum or maximum price was inside this range. Each firm's first step minimum quantity was ordinarily set to 0.8 times its Cournot quantity, and its maximum quantity was ordinarily set to 1.4 times its Cournot quantity, though occasionally these limits had to be widened by hand in order not to be binding.

In the second step, narrower ranges were set, so as to include all the non-zero-probability strategies found in the first step plus a margin equal to one first-stage grid step, in order to ensure that no non-zero probability would be overlooked. The second step solution was then checked to make sure that its limits were not binding, i.e. that the margin of the new 15×15 grid was all zeroes. Since *quadprog.m* only provides an approximate solution to the quadratic program, probabilities less than 10^{-6} were treated as zeroes for this purpose and in the graphs. The upper left and lower right corners of the second stage grids are plotted in the graphs as minute magenta or cyan dots.

Since occasionally solutions to the quadratic program were found that did not meet the additional zero-value restriction, *quadprog.m* was first initialized with the first strategies of both players. If the value of this program exceeded 10^{-4} in absolute value, the second strategies of both players were then tried, and so forth. Three different initializations were generally adequate.

Frequently *quadprog.m* returned error messages indicating that the problem was ill-conditioned, and occasionally it returned "solutions" which did not even come close to meeting the restrictions that were imposed. Often just changing the grid slightly gave much better results. Other algorithms exist for computing mixed Nash equilibria, and may give better conditioned computational results, but no Matlab programs for computing these turned up in a web search.