A NONPARAMETRIC ANALYSIS OF THE COURNOT MODEL

By Andres Carvajal and John K.-H. Quah

Abstract: An observer makes a finite number of observations of an industry producing a homogeneous good. Each observation consists of the market price, the output of individual firms and perhaps information on a firm's total cost or cost fluctuation. What restrictions on the data set are necessary and sufficient for it to be consistent with the hypothesis that firms in this industry are playing a Cournot game? We ask this question under two settings. In the first case, we assume that the observations are generated by changes in the market demand function, with each firm having a cost function that is stable across observations. We identify a set of inequalities that the observations must satisfy for it to be consistent with the Cournot hypothesis. In the second, we assume that the market demand function is stable across observations, but the data is generated by linear perturbations to the cost functions of firms in the industry. We show that the data is consistent with the Cournot model if and only if there exists a solution to a set of linear inequalities. Furthermore, the solution to these inequalities correspond (in some well-defined sense) to an estimate of the cost function of each firm.

Keywords: revealed preference, observable restrictions, linear programming, residual demand, Farkas' Lemma.

Authors' Emails: a.m.carvajal@warwaick.ac.uk, john.quah@economics.ox.ac.uk

Preliminary and Incomplete Manuscript. In particular, the Appendix (which will contain the proofs of many of the results) has yet to be written.

Comments welcome.

1. INTRODUCTION

Consider an industry with I firms producing a homogeneous good. We make T observations of this industry over time. How do we test the hypothesis that the firms in this industry are playing a Cournot game at each observation? We consider this problem under two settings.

In the first setting, we assume that the cost functions of firms in the industry do not vary across observations, so that the data is generated by fluctuations in the market demand function. The observation at t consists of the market price P_t, the output of the individual firm $Q_{i,t}$ (for every firm i), and the total cost incurred by the firm $C_{i,t}$. We identify the necessary and sufficient conditions that such a data set must satisfy for it to be consistent with the Cournot hypothesis. The main condition is a marginal property (M) that is easy to describe: the data set must not reveal instances of overproduction by firms. Specifically, suppose that at time t the firm is producing more than at some other time t', i.e., $Q_{i,t} > Q_{i,t'}$. Then the data must not reveal that the firm is better off at time t by reducing its output to $Q_{i,t'}$; specifically, condition M requires that

$$P_t Q_{i,t'} - C_{i,t'} < P_t Q_{i,t} - C_{i,t}.$$

The right hand side of this inequality is firm i's profit at time t if it produces $Q_{i,t}$. This is larger than the left hand side, which is an *under*-estimate of its profit at time t should it reduce output to $Q_{i,t'}$. Note that the left side is an under-estimate because the good's clearing price will be higher than P_t if firm i reduces its output (assuming, of course, that the market demand curve is downward sloping).

Not only is property M necessary for the Cournot hypothesis, we show that it is also, in essence, sufficient. This means that the observable restrictions on the Cournot model in this setting are rather weak. Indeed if the firms are colluding rather than playing a Cournot game, it will generate data that also satisfies M; in other words, collusion does not lead to any observable behavior that is inconsistent with the Cournot model. The main reason for this is that M is just a test that firms are not overproducing - it does not test for underproduction. This in turn follows from the fact that at each time t, the observer observes just one point of time t's market demand curve - so while it is possible that firms are curbing production at time t in a way that is not individually rational, it is also possible that each firm's output is individually rational (given the output of other firms), because the price will sharply below P_t should any firm increase production. In other words, there is too little information about the demand curve at time t for the observer to distinguish between these two cases.

The second part of this paper considers another setting in which the observable restrictions on the Cournot model are considerably stronger. In this case, we assume that the demand curve is stable across observations. Different observations over time are generated by perturbations to the cost functions of individual firms. The observation t consists of the price P_t, the output of each firm, $Q_{i,t}$, and the perturbation in marginal cost experienced by each firm, $\theta_{i,t}$. So the cost function of firm i consists of a permanent component $\bar{C}_i$ (which is not directly observable) and a perturbation $\theta_{i,t}$ which can be observed. Firm i's total cost of producing output q at time t is $C_{i,t}(q) = \bar{C}_i(q) + \theta_{i,t}q$. One could, for example, think of $\theta_{i,t}$ as a tax imposed on each unit of firm i's output that varies over time. In this setting, the problem is essentially one of constructing the unobserved permanent component of the cost functions $\bar{C}_i$ (for every firm i) in a way that makes the data consistent with the Cournot model.

In turns out that this problem has a solution if and only if there is a solution to a set of linear inequalities constructed from the data. In a specific sense, the solution to these inequalities also provide us with estimates of the permanent component $\bar{C}_i$ of firm i's cost function (for every firm i). The inequalities are constructed by revealed preference considerations, i.e., by requiring that firm i's choice at each time t is superior to all the other options which the data set has revealed to be open to firm i. Note that because demand is stable across observations, the market price and total output at some other time t' nonetheless provides the observer with information about firm i's residual demand curve at time t. Firm i's output choice at time t must be consistent with individually profit-maximizing behavior given this information.

Related Literature. This paper is a contribution to the literature on the observable restrictions/testable implications of various canonical economic models. One of the most influential papers in this literature is Afriat (1967), which identified the general axiom of revealed preference as the necessary and sufficient condition that a finite data set of price and demand observations must satisfy for it to be compatible with the utility-maximization hypothesis. This paper has generated a very large empirical literature. It has also been extended in various ways; in particular, see Varian (1982) for an extension to production theory and Brown and Matzkin (1996) for an analysis of observable restrictions in general equilibrium models. Our paper is also related to the large empirical IO literature (surveyed in Bresnahan (1989)) that aims either to test various IO models or to derive cost/demand information from observed behavior (typically under various game-theoretic assumptions).

Organization of the paper. The next section sets out some concepts used throughout the paper and also considers the problem of testing the profit-maximization hypothesis for a monopoly using a finite data set, under the assumption that cost is stable but demand is variable across observations. Section 3 considers an industry with several firms and characterizes Cournot rationalizability, again under the assumption that observations are generated by fluctuating demand. The case where demand is stable but observations are generated by cost perturbations is examined in Sections 4 and 5. Section 4 identifies necessary and sufficient conditions for a data set to be consistent with a profit-maximizing monopoly in this context; this leads naturally to a necessary and sufficient test for the Cournot hypothesis in Section 5.

2. RATIONALIZABILITY FOR A MONOPOLY WITH VARIABLE DEMAND

Consider an experiment in which we make T observations of a monopolist. The observations are indexed by t in the set $\mathcal{T} = \{1, 2, \ldots, T\}$; observation t consists of a triple (P_t, Q_t, Π_t), respectively the price charged by the monopolist, the quantity he sells, and the profit he makes. We require $P_t > 0$ and $Q_t > 0$ for all t; we also require the profit Π_t to be (strictly) smaller than total revenue P_tQ_t, so that the total cost C_t incurred by the monopolist in producing Q_t, which equals $P_tQ_t - \Pi_t$, is positive. The value of Π_t may be positive or negative, so we may observe losses.

We say that the set of observations $\{(P_t, Q_t, \Pi_t)\}_{t\in\mathcal{T}}$ is *rationalizable* if they are consistent with a profit-maximizing monopolist having a stable cost structure, with each observation corresponding to a different demand condition. Formally, we require that there be a $\mathbb{C}^1$ function $\bar{C} : R_+ \to R$ and $\mathbb{C}^1$ functions $\bar{P}_t : R_+ \to R$, for each t in $\mathcal{T}$, such that

(i) $\bar{C}(q) \geq 0$ and $\bar{C}'(q) > 0$;
(ii) $\bar{P}_t(q) \geq 0$ and $\bar{P}_t'(q) \leq 0$, with the the latter inequality being strict if $\bar{P}_t(q) > 0$;
(iii) $\bar{C}(Q_t) = C_t$ and $\bar{P}_t(Q_t) = P_t$; and
(iv) $\text{argmax}_{q \geq 0}[\bar{P}_t(q)q - \bar{C}(q)] = Q_t$.

Function $\bar{C}$ is the monopolist's *cost function*; condition (i) says that it is positive and strictly increasing.[1] Function $\bar{P}_t$ is the *inverse demand function* at observation t; condition (ii) says that more output can only be sold at a strictly lower price, until the price reaches zero. From this point on, we shall refer to any $\mathbb{C}^1$ cost function satisfying (i) as a *regular* cost function; similarly, a *regular* inverse demand function is a $\mathbb{C}^1$ inverse demand function that obeys (ii). Condition (iii) requires the inverse demand and cost functions to coincide with their observed values at each t. Lastly, condition (iv) requires the observations to be consistent with profit maximization. It is clear that conditions (iii) and (iv) together guarantee that the observed profit is $\Pi_t = \max_{q \geq 0}[\bar{P}_t(q)q - \bar{C}(q)]$. Note that we have allowed for the existence of sunk costs since we do not require $\bar{C}(0) = 0$. This implies that there is no nonnegativity constraint on profits, since the option of producing nothing can still incur a cost.

We say that the observations are *generic* if $Q_t \neq Q_{t'}$ whenever $t \neq t'$. Let $\{(P_t, Q_t, \Pi_t)\}_{t \in \mathcal{T}}$ be a generic set of observations. For each t, we define the set $S(t) = \{t' \in \mathcal{T} : Q_{t'} < Q_t\}$; in other words, $S(t)$ consists of those observations with output levels lower than Q_t. When $S(t)$ is nonempty, we denote $s(t) = \text{argmax}_{t' \in S(t)} Q_{t'}$; that is, $s(t)$ is the observation corresponding to the highest output level below Q_t. For those observations t with nonempty $S(t)$, we define $\Delta Q_t = Q_t - Q_{s(t)}$ and $\Delta C_t = C_t - C_{s(t)}$. So, ΔC_t is the extra cost incurred by the monopoly when it increases its output from $Q_{s(t)}$ to Q_t. We denote the *average marginal cost* over that output range by $M_t = \Delta C_t / \Delta Q_t$.

The generic set of observations $\{(P_t, Q_t, \Pi_t)\}_{t \in \mathcal{T}}$ is said to satisfy the *increasing cost property* (property C) if $\Delta C_t > 0$ whenever it is defined. It obeys the *marginal property* (property M) if, whenever $S(t)$ is nonempty,

$$P_t Q_{t'} - C_{t'} < P_t Q_t - C_t \quad \text{for } t' \in S(t). \tag{1}$$

---

[1] Note that while $C(0)$ may be zero or positive, (i) implies that $C(q) > 0$ if $q > 0$.

We may re-arrange this inequality to obtain

$$C_t - C_{t'} = \sum_{s \in S(t) \setminus (S(t') \cup \{t'\})} \Delta C_s < P_t(Q_t - Q_{t'}) \text{ for } t' \in S(t). \tag{2}$$

The observer does not know the exact market price of the product at time t, should the monopolist produce at $Q_{t'} < Q_t$, but he knows that it must be at least P_t. If Q_t is optimal, the cost saving in producing at $Q_{t'}$ rather than Q_t must be dominated by the revenue lost in producing at $Q_{t'}$ rather than Q_t; the latter does not exceed $P_t(Q_t - Q_{t'})$, so we obtain (2). In short, property M requires that the monopolist is not *overproducing* given the data.

PROPOSITION 1. *The generic set of observations* $\{(P_t, Q_t, \Pi_t)\}_{t \in \mathcal{T}}$ *is rationalizable only if it obeys properties C and M.*

Proof: If the set of observations is rationalizable, then for any t' in $S(t)$, we have $C_t - C_{t'} = \int_{Q_{t'}}^{Q_t} \bar{C}'(q)dq > 0$, since $C'(q) > 0$, so property C holds. Suppose that there is a violation of M. Then $P_t Q_{t'} - C_{t'} \geq P_t Q_t - C_t$ for t' in $S(t)$. But $\bar{P}_t(Q_{t'}) > \bar{P}_t(Q_t) = P_t$, so $\bar{P}_t(Q_{t'})Q_{t'} - C_{t'} > P_t Q_t - C_t$, which means that the monopolist is better off producing at $Q_{t'}$ rather than at Q_t. *Q.E.D.*

The next result says that properties C and M are also sufficient for rationalizability.

THEOREM 1. *Suppose the generic set of observations* $\{(P_t, Q_t, \Pi_t)\}_{t \in \mathcal{T}}$ *obeys C and M, and let* $\{\alpha_t\}_{t \in \mathcal{T}}$ *be a set of numbers satisfying* $0 < \alpha_t < P_t$. *Then the observations are rationalizable and the cost function* $\bar{C} : R_+ \to R$ *can be chosen such that* $\bar{C}'(Q_t) = \alpha_t$ *for all* $t \in \mathcal{T}$.

Note that the last condition in Theorem 1 says that we are free to choose the marginal cost at the optimal output level, subject to it being lower than the observed price. This feature will turn out to be useful in the next section, when we examine rationalizability in the Cournot model.

Since property M is a 'one-sided' condition - it requires that the monopoly is not *over*producing given the data - Theorem 1 effectively says that the data does not permit the observer to check that the monopolist is not *under*producing. The reason for this is that the monopolist decision at observation t to produce Q_t, *but not more*, can always be justified on the grounds that the price will fall sharply should it produce more. (This will be clearer with Lemmas 1 and 2 below.) The fact that the demand curve changes from one observation to the next, and the fact only one observation is made at each demand curve, means that such a possibility cannot be excluded.

Theorem 1 is an immediate consequence of the following two lemmas. Loosely speaking, Lemma 1 provides us with the cost function needed to rationalize the set of observations, while Lemma 2 gives the demand functions corresponding to each observation t.

LEMMA 1. *Suppose the generic set of observations* $\{(P_t, Q_t, \Pi_t)\}_{t \in \mathcal{T}}$ *obeys C and M and let* $\{\alpha_t\}_{t \in \mathcal{T}}$ *be a set of numbers satisfying* $0 < \alpha_t < P_t$. *Then, there is a regular cost function* $\bar{C} : R_+ \to R$ *such that, for all t in* $\mathcal{T}$,
(i) $\bar{C}(Q_t) = C_t$, *and* $\bar{C}'(Q_t) = \alpha_t$;
(ii) on a neighborhood of Q_t, $\bar{C}$ *is twice differentiable and satisfies that* $\bar{C}''(q) > 0$; *and*
(iii) for all q in $[0, Q_t)$,

$$P_t q - \bar{C}(q) < P_t Q_t - \bar{C}(Q_t). \tag{3}$$

Proof: See Appendix.

Property (i) in Lemma 1 requires the cost function to obey the specified marginal cost conditions and to agree with the cost data at the observed output levels. Property (ii) requires the cost function to be strictly concave in some neighborhood of the observed output levels. Finally, property (iii) is a strengthening of M: M requires (3) to hold at discrete output levels, while (ii) requires it to hold at *all* output levels up to Q_t.

The next result says that, for the cost function guaranteed by Lemma 1, we could find a demand function for each t such that the profit-maximizing output decision is Q_t.

LEMMA 2. *Let* $\{\alpha_t\}_{t \in \mathcal{T}}$ *be a set of numbers satisfying* $0 < \alpha_t < P_t$, *and let* $\bar{C} : R_+ \to R$ *be a regular cost function satisfying the three properties of Lemma 1. Then, for any* $t \in \mathcal{T}$, *there is a regular inverse demand function* $\bar{P}_t : R_+ \to R$ *such that*
(i) $\bar{P}_t(Q_t) = P_t$; *and*
(ii) $\text{argmax}_{q \geq 0} [\bar{P}_t(q) q - \bar{C}(q)] = Q_t$.

The proof of this lemma is in the Appendix, but the result is straightforward. The lemma requires that we produce an inverse demand function. Property (iii) in Lemma 1 already provides us with one such function that obeys (i) and (ii) (in Lemma 2): simply assume that $\bar{P}_t(q) = P_t$ for all $q \leq Q_t$ and $\bar{P}_t(q) = 0$ for $q > Q_t$. This function is non-regular, but we can always construct a regular demand function that is arbitrarily close to it.

3. COURNOT RATIONALIZABILITY WITH VARIABLE DEMAND

An industry consists of I firms producing a homogeneous good; we denote the set of firms by $\mathcal{I} = \{1, 2, \ldots, I\}$. Consider an experiment in which T observations are made of this industry. As in the previous section, we index the observations by t in $\mathcal{T} = \{1, 2, \ldots, T\}$. For each t, the industry price P_t, the output of each firm $(Q_{i,t})_{i\in\mathcal{I}}$ and their profits $(\Pi_{i,t})_{i\in\mathcal{I}}$ are observed. We require $P_t > 0$ and $Q_{i,t} > 0$ for all t and i; the profit observations $\Pi_{i,t}$ can take either positive and negative values. Note that the total cost incurred by firm i in producing $Q_{i,t}$, which we denote by $C_{i,t}$, follows immediately from the equation $C_{i,t} = P_t Q_{i,t} - \Pi_{i,t}$. For simplicity of notation, we henceforth let $Q_t = \sum_{i\in\mathcal{I}} Q_{i,t}$ denote the aggregate output of the industry at observation t.

We say that the set of observations $\{[P_t, (Q_{i,t})_{i\in\mathcal{I}}, (\Pi_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ is *Cournot rationalizable* if each observation can be explained as a Cournot equilibrium arising from a different market demand function, keeping the cost function of each firm fixed across observations. Formally, we require that there be a regular cost function $\bar{C}_i : R_+ \to R$ for each firm i and a regular demand function $\bar{P}_t : R_+ \to R$ for each t, such that

(i) $\bar{C}_i(Q_{i,t}) = C_{i,t}$ and $\bar{P}_t(Q_t) = P_t$; and

(ii) $\text{argmax}_{q_i \geq 0}[q_i \bar{P}_t(q_i + \sum_{j\neq i} Q_{j,t}) - \bar{C}_i(q_i)] = Q_{i,t}$.

Condition (i) says that the inverse demand and cost functions must coincide with their observed values at each t. Condition (ii) says that, at each observation t, firm i's observed output level $Q_{i,t}$ maximizes its profit given the output of the other firms. It is clear that these conditions imply that the observed profit $\Pi_t = \max_{q_i\geq 0}[q_i \bar{P}_t(q_i + \sum_{j\neq i} Q_{j,t}) - \bar{C}_i(q_i)]$. Note that, as in the previous section, we allow for the existence of sunk costs, since we do not require $\bar{C}_i(0) = 0$.

We say that the observations are *generic* if, for every firm i, we have $Q_{i,t} \neq Q_{i,t'}$ whenever $t \neq t'$. Let $\{[P_t, (Q_{i,t})_{i\in\mathcal{I}}, (\Pi_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ be a generic set of observations. For each firm

i, we may define $S_{i,t}$, $s_{i,t}$, $\Delta Q_{i,t}$, $\Delta C_{i,t}$, and $M_{i,t}$, in a way analogous to our definitions in the previous section. We say that the set of observations obey the increasing cost property if, for each i, $\{(P_t, Q_{i,t}, \Pi_{i,t})\}_{t\in\mathcal{T}}$ obeys property C (in the sense previously defined). Similarly, we say that it obeys the marginal property if, for each i, $\{(P_t, Q_{i,t}, \Pi_{i,t})\}_{t\in\mathcal{T}}$ obeys property M.

It is clear, for exactly the same reasons as the ones given in the monopoly case, that properties C and M are *necessary* for a set of observations to be Cournot rationalizable. Specifically, C is needed to guarantee that each firm's production cost is increasing in output, and M is needed to guarantee that each firm is not strictly better off by producing less than the observed output. The next result says that these conditions are also sufficient for Cournot rationalizability.

THEOREM 2. *Suppose that the generic set of observations* $\{[P_t, (Q_{i,t})_{i\in\mathcal{I}}, (\Pi_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ *obeys properties C and M. Then the set is Cournot rationalizable.*

Just as Theorem 1 follows from Lemmas 1 and 2, we can prove Theorem 2 with a similar two-step procedure. Note that, at observation t, if firm i is indeed playing its best response for demand function $\bar{P}_t$ and cost function $\bar{C}_i$, then the first order condition

$$\bar{P}_t'(Q_t)Q_{i,t} + P_t = \bar{C}_i'(Q_{i,t}) \tag{4}$$

must be satisfied. It follows that

$$-\bar{P}_t'(Q_t) = \frac{P_t - \bar{C}_1'(Q_{1,t})}{Q_{1,t}} = \frac{P_t - \bar{C}_2'(Q_{2,t})}{Q_{2,t}} = \ldots = \frac{P_t - \bar{C}_I'(Q_{I,t})}{Q_{I,t}}. \tag{5}$$

This motivates the condition imposed on the cost functions in the next result, which is loosely analogous to Lemma 1.

LEMMA 3. *Let* $\{[P_t, (Q_{i,t})_{i\in\mathcal{I}}, (\Pi_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ *be a generic set of observations obeying C and M and suppose that the positive numbers* $\{\alpha_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ *satisfy*

$$\frac{P_t - \alpha_{1,t}}{Q_{1,t}} = \frac{P_t - \alpha_{2,t}}{Q_{2,t}} = \ldots = \frac{P_t - \alpha_{I,t}}{Q_{I,t}} > 0 \text{ for all } t \text{ in } \mathcal{T}. \tag{6}$$

Then, there are regular cost functions $\bar{C}_i : R_+ \to R$ *such that*
(i) $\bar{C}_i(Q_{i,t}) = C_{i,t}$ *and* $\bar{C}_i'(Q_{i,t}) = \alpha_{i,t}$*;*

(ii) on a neighborhood of $Q_{i,t}$, $\bar{C}_i$ is twice differentiable and satisfies that $\bar{C}_i''(q) > 0$; and
(iii) for all q_i in $[0, Q_{i,t})$,

$$P_t q_i - \bar{C}_i(q_i) < P_t Q_{i,t} - \bar{C}_i(Q_{i,t}). \tag{7}$$

Proof: See Appendix.

It is important to notice that for *any* P_t and $\{Q_{i,t}\}_{i\in\mathcal{I}}$ there *always* exist positive numbers $\{\alpha_{i,t}\}_{i\in\mathcal{I}}$ such that equation (6) holds. Suppose that firm k produces more than any other firm at observation t, i.e., $Q_{k,t} \geq Q_{i,t}$ for all i in I. Let $\alpha_{k,t}$ be any positive number smaller than P_t, and define $\beta = (P_t - \alpha_{k,t})/Q_{k,t}$. Then,

$$\alpha_{i,t} = P_t - \beta Q_{i,t} \geq P_t - \beta Q_{k,t} = \alpha_{k,t} > 0.$$

The next result is analogous to Lemma 2. It is clear that this result together with Lemma 3 proves Theorem 2.

LEMMA 4. *Let $\{\alpha_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ be a set of positive numbers satisfying equation (6) and suppose that the cost functions $\bar{C}_i : R_+ \to R$ satisfy the properties in Lemma 3. Then, there are regular demand functions $\bar{P}_t : R_+ \to R$ such that, for every firm i,*

$$\text{argmax}_{q_i\geq 0}[q_i \bar{P}_t(q_i + \sum_{j\neq i} Q_{j,t}) - \bar{C}_i(q_i)] = Q_{i,t}.$$

Like Theorem 1, Theorem 2 has the feature that it only checks that each firm in the Cournot oligopoly is not overproducing - in this setting, the data cannot effectively test there is no underproduction. Properties C and M are sufficiently weak that there are other reasonable scenarios of firm interaction under which they will also be satisfied. In particular, these properties hold if the firms are colluding to maximize total profits, which means that their collusion will never generate any evidence that is contrary to the Cournot model.

To state this claim formally, we define a dataset $\{[P_t, (Q_{i,t})_{i\in\mathcal{I}}, (\Pi_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ as being *consistent with collusion* if there is a regular cost function $\bar{C}_i : R_+ \to R$ for each firm i and a regular demand function $\bar{P}_t : R_+ \to R$ for each t, such that
(i) $\bar{C}_i(Q_{i,t}) = C_{i,t}$ and $\bar{P}_t(Q_t) = P_t$; and
(ii) $\text{argmax}_{(q_i,t)_{i\in\mathcal{I}}\geq 0} \left[\sum_{i\in\mathcal{I}} \bar{P}_t(\sum_{i\in\mathcal{I}} q_{i,t}) - \sum_{i\in\mathcal{I}} \bar{C}_i(q_i)\right] = (Q_{i,t})_{i\in\mathcal{I}}$.
Condition (i) requires that the cost and inverse demand functions agree with the data while (ii) says that the observed output across firms must maximize total profit.

Corollary 1. *The dataset* $\{[P_t, (Q_{i,t})_{i\in\mathcal{I}}, (\Pi_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ *is consistent with collusion only if it obeys C and M.*

Proof: As usual, C is needed for the cost function of each firm to be strictly increasing. Suppose that M is violated, so for some firm i', observation t, and $t' \in S_{i'}(t)$,

$$P_t Q_{i',t'} - C_{i',t'} \geq P_t Q_{i',t} - C_{i',t} = \Pi_{i',t}.$$

By definition, $Q_{i',t'} < Q_{i',t}$, so $\bar{P}_t(Q_{i',t'} + \sum_{j\neq i'} Q_{j,t}) > P_t$. This implies that

$$\bar{P}_t\left(Q_{i',t'} + \sum_{j\neq i'} Q_{j,t}\right) Q_{i',t'} - C_{i'}(Q_{i',t'}) > \Pi_{i',t}, \tag{8}$$

whereas for every $i \neq i'$,

$$\bar{P}_t\left(Q_{i',t'} + \sum_{j\neq i'} Q_{j,t}\right) Q_{i,t} - C_i(Q_{i,t}) > \Pi_{i,t}. \tag{9}$$

In other words, both i' and all other firms are strictly better off if i' reduces its output (with the other firms benefitting from the higher market price). Clearly, the output vector $(Q_{i,t})_{i\in\mathcal{I}}$ does not maximize total industry profit at t, so collusion is excluded. *Q.E.D.*

Theorem 2 says that, in a variable demand environment, there *are* observable restrictions to the Cournot model, albeit rather weak restrictions. We now consider the same problem, but assume that the profits earned - and thus the costs incurred - by each firm are not known. Formally, the data set reduces to $\{(P_t, (Q_{i,t})_{i\in\mathcal{I}})\}_{t\in\mathcal{T}}$. This data set is said to be *generic* if $Q_{i,t'} \neq Q_{i,t}$ whenever $t \neq t'$; it is *Cournot rationalizable* if we can find a regular demand function, $\bar{P}_t$, for each observation t, and a regular cost function, $\bar{C}_i$, for each firm i, such that
(i) $\bar{P}_t(\sum_{i\in\mathcal{I}} Q_{i,t}) = P_t$; and
(ii) $\text{argmax}_{q_i\geq 0}[q_i \bar{P}_t(q_i + \sum_{j\neq i} Q_{j,t}) - \bar{C}_i(Q_i)] = Q_{i,t}$.
In words, the tth observation, $(P_t, (Q_{i,t})_{i\in I})$, is the Cournot outcome when each firm i has cost function $\bar{C}_i$ and the market inverse demand function is $\bar{P}_t$. The following result says that Cournot competition imposes no restriction on the observations $\{(P_t, (Q_{i,t})_{i\in\mathcal{I}})\}_{t\in\mathcal{T}}$.

Corollary 2. *Any generic set of observations* $\{(P_t, (Q_{i,t})_{i\in\mathcal{I}})\}_{t\in\mathcal{T}}$*, is Cournot rationalizable.*

Proof: By Theorem 2, it suffices that we find an array of individual costs, $\{C_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$, that, when added to the observed data, gives a set of observations that obeys C and M. To see that this is possible, let $\mu_i = \min_{t:S(t)\neq\varnothing}\{P_t(Q_{i,t} - Q_{i,s_i(t)})\}$. Since $\mu_i > 0$, we can pick $\{C_{i,t}\}_{t\in\mathcal{T}}$ satisfying C and with $0 < C_{i,t} < \mu_i$ for every t. Whenever $Q_{i,t'} < Q_{i,t}$,

$$C_{i,t} + P_t(Q_{i,t'} - Q_{i,t}) \leq C_{i,t} + P_t(Q_{i,s_i(t)} - Q_{i,t}) \leq C_{i,t} - \mu_i < 0 < C_{i,t'}$$

so $\{C_{i,t}\}_{t\in\mathcal{T}}$ also obeys M. *Q.E.D.*

Corollary 2 holds because, in a sense, we are free to choose the cost functions. If we restrict these functions then there *will* be restrictions on $\{(P_t, (Q_{i,t})_{i\in\mathcal{I}})\}_{t\in\mathcal{T}}$. For example, it is easy to show that at any Cournot equilibrium, firms' market shares vary inversely with their marginal costs. This means that if firms have constant marginal costs (i.e., constant across each firm's output level), then the *ranking* of firms by market shares will not vary across observations, even as demand fluctuates. Clearly, this is restriction on the data set.

4. Rationalizability for a monopoly with cost perturbations

We return once again to the rationalizability problem for a monopoly. Unlike Section 2, we now consider a setting where the demand function is stable across observations, but there are observed perturbations to the monopolist's cost function. These perturbations take the form of a change in marginal cost that is constant at all output levels. For example, a unit tax imposed on the monopolist's output that is different at each observation will constitute such a perturbation.

An observation (at t) consists of the price charged by the monopolist P_t, its output Q_t, and the change in marginal cost θ_t. We assume that P_t and Q_t are strictly positive; without loss of generality, we also assume that $\theta_t \geq 0$.[2] Furthermore, we require that $P_t > \theta_t$ (otherwise, the firm will not choose to produce anything) and also that $P_t > P_{t'}$ whenever $Q_t < Q_{t'}$ (so the monopolist faces a downward sloping demand curve).

We say that the data set $\{(P_t, Q_t, \theta_t)\}_{t\in\mathcal{T}}$ is *generic* if $Q_{i,t} \neq Q_{i,t'}$ whenever $t \neq t'$. This data set is *rationalizable* if it is consistent with the behavior of a profit-maximizing

---

[2] The lowest value of θ_t (amongst all the observations) can always be normalized at zero.

monopolist. Formally, we require a $\mathbb{C}^1$ function $\bar{P} : R_+ \to R$ and a $\mathbb{C}^1$ function $\bar{C} : R_+ \to R$ such that
(i) $\bar{C}(q) \geq 0$ and $\bar{C}'(q) > 0$;
(ii) $\bar{P}(q) \geq 0$ and $\bar{P}'(q) \leq 0$, with the latter inequality being strict when $\bar{P}(q) > 0$;
(iii) $\bar{P}(Q_t) = P_t$; and
(iv) $\text{argmax}_{q \geq 0}[q\bar{P}(q) - \bar{C}(q) - \theta_t q] = Q_t$.
Note that $\bar{P}$ should be interpreted as the inverse demand function faced by the monopolist, while the cost function of the monopolist at observation t is C_t, with $C_t(q) = \bar{C}(q) + \theta_t q$. This cost function has two parts: a *permanent* component $\bar{C}$ which does not change across observations and a perturbation θ_t.

We wish to identify the restrictions that a rationalizable data set $\{P_t, (Q_t, \theta_t)\}_{t \in \mathcal{T}}$ must obey. Before we consider that, it is useful to introduce some notation. Recall that $s(t)$ refers to that observation whose output is just below Q_t. We denote the cost of increasing output from $Q_{s(t)}$ to Q_t by c_t, i.e., $c_t = \bar{C}(Q_t) - \bar{C}(Q_{s(t)})$. (So this is the increase in the permanent component of the cost; it excludes the perturbation θ_t.) Abusing notation somewhat, we denote the set $\{s : Q_{t'} < Q_s \leq Q_t\}$ by $(t', t]$. Therefore, $\sum_{s \in (t',t]} c_s$ is the cost of increasing output from Q_t' to Q_t, i.e., $\sum_{s \in (t',t]} c_s = \bar{C}(Q_t) - \bar{C}(Q_{t'})$.

Consider two observations with $Q_{t'} < Q_t$. At observation t, the firm's added revenue from increasing production from $Q_{t'}$ to Q_t (net of the marginal cost θ_t) must exceed the added permanent cost of this increase in output. In other words,

$$P_t Q_t - P_{t'} Q_{t'} - \theta_t (Q_t - Q_{t'}) > \sum_{s \in (t',t]} c_s. \tag{10}$$

On the other hand, at observation t', the firm's added revenue from increasing production from $Q_{t'}$ to Q_t must be exceeded by the cost of such an increase. This means that

$$P_t Q_t - P_{t'} Q_{t'} - \theta_{t'} (Q_t - Q_{t'}) < \sum_{s \in (t',t]} c_s. \tag{11}$$

The next result says that restrictions of the type (10) and (11) are not just necessary for rationalizability - they are also sufficient.

THEOREM 3. *The generic data set $\{P_t, Q_t, \theta_t\}_{t \in \mathcal{T}}$ is rationalizable if and only if there are numbers $c_t > 0$ (for $t \in \mathcal{T}$) that obey all inequalities of the type (10) and (11).*

In the case where the data set is rationalizable, the function $\bar{C}$ *can be chosen such that* $\bar{C}(Q_t) - \bar{C}(Q_{s(t)}) = c_t$.

Therefore, just as in Afriat's Theorem, determining rationalizability in this context is equivalent to finding a solution to a set of linear inequalities. These linear inequalities take a simple form. For example, suppose that $T = 4$ and (to keep the notation simple) assume that $Q_1 < Q_2 < Q_3 < Q_4$. Then the conditions (10) and (11) require that there be positive numbers c_2, c_3 and c_4 that satisfy the inequalities

$$a^1 < c_2 < A^1; \quad a^2 < c_3 < A^2; \quad a^3 < c_4 < A^3$$

$$a^4 < c_2 + c_3 < A^4; \quad a^5 < c_3 + c_4 < A^5$$

$$a^6 < c_2 + c_3 + c_4 < A^6,$$

where the upper and lower bounds (A^m and a^m) are numbers given by the left side of (10) and (11) respectively. This pattern of inequalities is maintained whatever the number of observations. A simple calculation shows that for a data set with T observations, one is required to find positive constants c_2, c_3,..., c_T that satisfy $T(T-1)$ linear inequalities.

Theorem 4 gives a simple, easily implementable, way of checking rationalizability, but the conditions may seem a little opaque. There is an alternative - and very intuitive - characterization of rationalizability that we shall now describe.

Consider the observations as occurring in T different states of the world, with t occurring with some strictly positive probability π_t. A strategy $a = (a_t)_{t \in \mathcal{T}}$ for the monopolist specifies an action a_t for each realized value of t; an action is a lottery over the set $\{Q_t : t \in \mathcal{T}\}$. A lottery that, at state t, specifies Q_s with probability ρ_s has the expected revenue (conditional on state t) of $\sum_{s \in \mathcal{T}} \rho_s (P_s - \theta_t) Q_s$; we denote this by $R(a_t)$. We denote the ex ante distribution over $\{Q_t : t \in \mathcal{T}\}$ induced by $\pi = (\pi_t)_{t \in \mathcal{T}}$ and the strategy a by $D(\pi, a)$.

We may interpret the data set $\{P_t, Q_t, \theta_t\}_{t \in \mathcal{T}}$ as revealing that the monopolist has chosen the strategy $\hat{Q} = \{Q_t\}_{t \in \mathcal{T}}$. We say that the strategy $\hat{Q}$ is *irrational in expectation at* π if there is an alternative strategy a such that

(i) $\sum_{t \in \mathcal{T}} \pi_t R(a_t) > \sum_{t \in \mathcal{T}} \pi_t R(Q_t)$ and

(ii) $D(\pi, \hat{Q})$ first order stochastically dominates $D(\pi, a)$.

Note that (ii) implies that the ex ante average cost of strategy a is lower than that of $\hat{Q}$, for any increasing cost function. Therefore, (i) and (ii) imply that the ex ante average profit of $\hat{Q}$ must be lower than a, which means that $\hat{Q}$ cannot be ex post optimal (i.e., with Q_t optimal at every state $t \in \mathcal{T}$) for *any* cost function. We conclude that if there is some π at which $\hat{Q}$ is irrational in expectation, then the data set $\{P_t, Q_t, \theta_t\}_{t \in \mathcal{T}}$ is not rationalizable. The next result says that the converse is also true.

THEOREM 4. *The data set* $\{P_t, Q_t, \theta_t\}_{t \in \mathcal{T}}$ *is rationalizable if and only if there does not exist* π *for which the strategy* $\hat{Q} = (Q_t)_{t \in \mathcal{T}}$ *is irrational in expectation.*

The proof of this result uses Farkas' Lemma and can be found in in the Appendix. That it uses Farkas' Lemma should not be surprising; indeed this result has the flavor of the no-arbitrage theorem: loosely speaking, rationalizability is analogous to the existence of positive state prices for a set of asset prices, and the absence of π at which $\hat{Q}$ is irrational in expectation is analogous to the absence of arbitrage.

5. COURNOT RATIONALIZABILITY WITH COST PERTURBATIONS

Consider an industry with I firms, where industry demand (for the homogeneous good) is stable across observations. At each observation, there is a perturbation to the cost function of each firm that is observable. The observation t consists of the price P_t, the output of each firm, $Q_{i,t}$, and the perturbation in marginal cost experienced by each firm, $\theta_{i,t}$. We make assumptions analogous to those made in the previous section: P_t and Q_t are strictly positive, $\theta_{i,t}$ is nonnegative, $P_t > \theta_{i,t}$, and $P_t > P_{t'}$ whenever $Q_t < Q_{t'}$.

The data set $\{P_t, (Q_{i,t})_{i \in \mathcal{I}}, (\theta_{i,t})_{i \in \mathcal{I}}\}_{t \in \mathcal{T}}$ is said to be *generic* if $Q_{i,t} \neq Q_{i,t'}$ whenever $t \neq t'$. It is *Cournot rationalizable* if there is a $\mathbb{C}^1$ function $\bar{P} : R_+ \to R$ and a $\mathbb{C}^1$ function $\bar{C}_i : R_+ \to R$ for each firm i in I, such that

(i) $\bar{C}_i(q) \geq 0$ and $\bar{C}_i'(q) > 0$;

(ii) $\bar{P}(q) \geq 0$ and $\bar{P}'(q) \leq 0$, with the latter inequality being strict when $\bar{P}(q) > 0$;

(iii) $\bar{P}(Q_t) = P_t$; and

(iv) $\text{argmax}_{q_i \geq 0}[q_i \bar{P}(q_i + \sum_{j \neq i} Q_{j,t}) - (\bar{C}_i(q_i) + \theta_{i,t} q_i)] = Q_{i,t}$.

These conditions say that $\{Q_{i,t}\}_{i \in \mathcal{I}}$ can be thought of as the Cournot equilibrium at observation t, where $\bar{P}$ is interpreted as the inverse demand function and the cost function of firm i at observation t is $C_{i,t}$, with $C_{i,t}(q) = \bar{C}_i(q) + \theta_{i,t}q$, with $\bar{C}_i$ being the permanent component of firm i's cost function.

We wish to consider the restrictions that a Cournot rationalizable data set $\{P_t, (Q_{i,t})_{i\in\mathcal{I}}, (\theta_{i,t})_{i\in\mathcal{I}}\}_{t\in\mathcal{T}}$ must obey. This analysis is more complicated (and certainly heavier in notation) than the case for a monopolist considered in the previous section but it is not fundamentally different. The basic idea is to make use of the information on demand conditions drawn from the *entire* data set to impose restrictions on firm i's residual demand curve at a typical observation t. To be specific, let $\tilde{Q}_{i,t,t'}$ be the output level with which at observation t, firm i induces the aggregate output observed at t', namely $\tilde{Q}_{i,t,t'} = Q_{t'} - \sum_{j\neq i} Q_{j,t}$. The set $\{\tilde{Q}_{i,t,t'}, P_{t'}\}_{t'\in\mathcal{T}}$ consists of points on firm i's residual demand curve at observation t. The information provided by these points allows us to impose four types of restrictions on the data set which.

(i) Consider firstly the case where $\tilde{Q}_{i,t,t'} < Q_{i,t}$. Since the output $\tilde{Q}_{i,t,t'}$ was not chosen, the added revenue in increasing production to $Q_{i,t}$ must have exceeded the added cost. This implies that

$$P_tQ_{i,t} - P_{t'}\tilde{Q}_{i,t,t'} - \theta_{i,t}(Q_{i,t} - \tilde{Q}_{i,t,t'}) > \sum_{s\in\tilde{\mathcal{S}}_d(i,t,t')} c_{i,s}, \tag{12}$$

where

$$\tilde{\mathcal{S}}_d(i,t,t') = \{s \in T : \min_{s'}\{Q_{i,s'} : Q_{i,s'} \geq \tilde{Q}_{i,t,t'}\} < Q_{i,s} \leq Q_{i,t}\}.$$

The left of (12) gives the added revenue (net of the marginal cost $\theta_{i,t}$) from increasing production from $Q_{i,t,t'}$ to $Q_{i,t}$. This is greater than the (permanent) cost of increasing production from $Q_{i,t,t'}$ to $Q_{i,t}$, which in turn is bounded below by the expression on the right of (12).[3]

(ii) When $\tilde{Q}_{i,t,t'} > Q_{i,t}$, the output $\tilde{Q}_{i,t,t'}$ was not chosen because the added revenue of this increases production is dominated by the increase in cost. This implies that

$$P_{t'}\tilde{Q}_{i,t,t'} - P_tQ_{i,t} - \theta_{i,t}(\tilde{Q}_{i,t,t'} - Q_{i,t}) < \sum_{s\in\tilde{\mathcal{S}}_u(i,t,t')} c_{i,s}, \tag{13}$$

where

$$\tilde{\mathcal{S}}_u(i,t,t') = \{s \in T : Q_{i,t} < Q_{i,s} \leq \min_{s'}\{Q_{i,s'} : Q_{i,s'} \geq \tilde{Q}_{i,t,t'}\}\}.$$

---

[3] We are adapting the notation of the previous section in the obvious way, so $c_{i,t} = \bar{C}(Q_{i,t}) - C(Q_{s_i(t)})$, where $Q_{s_i(t)}$ is an observed output level just below Q_t.

Note that the right side of (13) is an upper bound on the cost of increasing output from $Q_{i,t}$ to $Q_{i,t,t'}$.

(iii) Define $T(i,t,\hat{t}) = \text{argmin}_s\{\tilde{Q}_{i,t,s} : \tilde{Q}_{i,t,s} \geq Q_{i,\hat{t}}\}$. Consider $\hat{t}$ such that $Q_{i,\hat{t}} < Q_{i,t}$; since $Q_{i,t}$ was chosen over $Q_{i,\hat{t}}$ at observation t, the increase in revenue from increasing production from $Q_{i,\hat{t}}$ to $Q_{i,t}$ must exceed the increase in cost. This implies that

$$P_tQ_{i,t} - P_{T(i,t,\hat{t})}Q_{i,\hat{t}} - P_tQ_{i,t} - \theta_{i,t}(Q_{i,t} - Q_{i,\hat{t}}) > \sum_{s\in\mathcal{S}_d(i,t,\hat{t})} c_{i,s}, \tag{14}$$

where

$$\mathcal{S}_d(i,t,\hat{t}) = \{s \in T : \min_{s'}\{Q_{i,s'} : Q_{i,s'} > Q_{i,\hat{t}}\} \leq Q_{i,s} \leq Q_{i,t}\}.$$

The left of (14) is is an upper bound of the increased revenue from raising production to $Q_{i,t}$ from $Q_{i,\hat{t}}$ since the price at output $Q_{i,\hat{t}}$ exceeds $P_{T(i,t,\hat{t})}$. This value is greater than the cost of raising output from $Q_{i,\hat{t}}$ to $Q_{i,t}$, which is the expression on the right of (14).

(iv) Lastly, consider $\hat{t}$ such that $Q_{i,\hat{t}} > Q_{i,t}$; given that $Q_{i,t}$ was chosen over $Q_{i,\hat{t}}$ at observation t, the increase in revenue from increasing production from $Q_{i,t}$ to $Q_{i,\hat{t}}$ must be less than the increase in cost. This implies that

$$P_{T(i,t,\hat{t})}Q_{i,\hat{t}} - P_tQ_{i,t} - \theta_{i,t}(Q_{i,\hat{t}} - Q_{i,t}) < \sum_{s\in\mathcal{S}_u(i,t,\hat{t})} c_{i,s}, \tag{15}$$

where

$$\mathcal{S}_u(i,t,\hat{t}) = \{s \in T : \min_{s'}\{Q_{i,s'} : Q_{i,s'} > Q_{i,t}\} \leq Q_{i,s} \leq Q_{i,\hat{t}}\}.$$

Note that the left of this inequality is an lower bound of the increased revenue from raising production to $Q_{i,\hat{t}}$ from $Q_{i,t}$ since the price at output $Q_{i,\hat{t}}$ exceeds $P_{T(i,t,\hat{t})}$.

The next result says that the restrictions (12) to (15) are not just necessary for a data set to be Cournot rationalizable - they are also sufficient. The proof of the sufficiency of these conditions is found in the Appendix. Note that as in Theorem 4, the conditions take the form of a set of linear inequalities.

THEOREM 5. *The generic data set* $\{P_t, (Q_{i,t})_{i\in\mathcal{I}}, (\theta_{i,t})_{i\in\mathcal{I}}\}_{t\in\mathcal{T}}$ *is Cournot rationalizable if and only if for each firm* i*, there are numbers* $c_{i,t} > 0$ *(for* $t \in \mathcal{T}$*) that obey the inequalities (12) to (15). In the case where the data set is rationalizable, the function* $\bar{C}_i$ *can be chosen such that* $\bar{C}_i(Q_{i,t}) - \bar{C}(Q_{s_i(t)}) = c_{i,t}$.

APPENDIX (to be written)

REFERENCES (incomplete)

AFRIAT, S (1967): The construction of a utility function from expenditure data, *International Economic Review,* 8, 67-77.

BROWN, D. AND R. MATZKIN (1996): Testable restrictions on the equilibrium manifold, *Econometrica*, 64-6, 1249-1262.

BRESNAHAN, T. F. (1989): Empirical Studies of Industries with Market Power, in *Hanbook of Industrial Economics, Vol. II*, eds. R.SCHMALENSEE AND R. D. WILLIG, Elsevier.

VARIAN, H (1982): The nonparametric approach to production analysis, *Econometrica*, 52, 579-597.