ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 81, No. 6 (November, 2013), 2351–2379

REVEALED PREFERENCE TESTS OF THE COURNOT MODEL

ANDRÉS CARVAJAL
University of Western Ontario, London, Ontario, N6A 5C2, Canada

RAHUL DEB
University of Toronto, Toronto, Ontario, M5S 3G7, Canada

JAMES FENSKE
Oxford University, Oxford, OX1 3UQ, United Kingdom

JOHN K.-H. QUAH
Oxford University, Oxford, OX1 3UQ, United Kingdom

Econometrica, Vol. 81, No. 6 (November, 2013), 2351–2379

REVEALED PREFERENCE TESTS OF THE COURNOT MODEL

BY ANDRÉS CARVAJAL, RAHUL DEB, JAMES FENSKE, AND JOHN K.-H. QUAH[1]

The aim of this paper is to develop revealed preference tests for Cournot equilibrium. The tests are akin to the widely used revealed preference tests for consumption, but have to take into account the presence of strategic interaction in a game-theoretic setting. The tests take the form of linear programs, the solutions to which also allow us to recover cost information on the firms. To check that these nonparametric tests are sufficiently discriminating to reject real data, we apply them to the market for crude oil.

KEYWORDS: Cournot equilibrium, nonparametric test, observable restrictions, linear programming, collusion, crude oil market.

1. INTRODUCTION

IN AN INFLUENTIAL PAPER, Afriat (1967) identified necessary and sufficient conditions for a finite set of observations of price vectors and demand bundles to be consistent with utility-maximizing behavior. These conditions take the form of a linear program that can be easily solved and, partly for this reason, a large literature on consumer behavior has been built on Afriat's result. A natural extension of Afriat's theorem is to derive observable restrictions on outcomes in a general equilibrium setting; this was carried out by Brown and Matzkin (1996), who gave a revealed preference analysis of Walrasian equilibria in an exchange economy.

The problem studied by Brown and Matzkin can be posed in other multiagent contexts. This paper raises a similar question in the context of an oligopoly. The basic model introduced in Section 2 considers a researcher who has access to a set of observations of an industry that produces a single good; each observation consists of the price of the good and the output of each firm. We ask whether there are any observable restrictions (i.e., restrictions on the data set) implied by the *Cournot hypothesis*: that each observation of price and output quantities constitutes a (static) Cournot equilibrium, with firms having

[1]The authors would like to thank a co-editor and three referees for their very helpful comments. The financial support of the ESRC to this research project, through Grants RES-000-22-3771 (Andrés Carvajal) and RES-000-22-3187 (John Quah) is gratefully acknowledged. Rahul Deb and James Fenske would like to acknowledge the financial support they received from their Leylan fellowships. Part of this paper was written while John Quah was visiting the Economics Department at the National University of Singapore and while Carvajal was visiting CORE; they thank these institutions for their hospitality. We are grateful for helpful comments from Dirk Bergemann, Don Brown, Greg Crawford, Luis Corchon, Françoise Forges, Enrico Minelli, Kirill Pogorelskiy, Margaret Slade, Mike Waterson, Mike Whinston, and Glen Weyl and seminar participants at SAET 2010 (Singapore), EWGET 2010 (Cracow), Workshop on Revealed Preference 2010 (Paris, Dauphine), LSE, Oxford, Washington University, Wisconsin-Madison, ASU, Vigo, and Tor Vergata.

 DOI: 10.3982/ECTA10175

cost functions that do not vary across observations and with the data generated by changes to a downward-sloping demand function. In keeping with the revealed preference approach, we make no parametric assumptions about demand and cost functions, and the object is to develop tests that are the most powerful possible given the data (in the sense of being both necessary and sufficient). Notice also that the observed data are extremely parsimonious; in particular, we do not assume that the observer has any knowledge of variables that are known to shift or twist the demand curves.

We show that the Cournot hypothesis as stated imposes no observable restrictions, that is, any data set is consistent with the Cournot hypothesis. However, we demonstrate through examples that there *are* observable restrictions imposed by this model as long as we require firms' cost curves to be convex (in other words, that they have increasing marginal costs). Convexity of the cost functions is a common assumption to make in industrial organization (IO) models since it helps to guarantee that each firm faces a quasiconcave profit-maximization problem. We go on to show that a data set is Cournot rationalizable with convex cost functions if and only if there is a solution to a linear program that we can explicitly construct from the data. So this gives an easy-to-implement and nonvacuous test of Cournot rationalizability.[2] This test is useful even in situations where the output of one or more firms is missing from the data set. This is because whenever some set of firms in an industry is playing a Cournot game, then any subset of those firms will also be playing a Cournot game.

If a data set is Cournot rationalizable with convex cost functions, there will be a family of convex cost functions (with one cost function for each firm) that can account for each observation as a Cournot equilibrium under some demand curve. Typically, there will be many such families; we show in Section 3 that this collection of convex cost function families is isomorphic (in some sense) to a compact and convex sublattice of the Euclidean space. Any compact and convex sublattice contains its supremum and its infimum; the former is associated with the family of cost functions with the highest marginal costs, while the latter is associated with the family of cost functions with the lowest marginal costs. We provide very intuitive algorithms for calculating the supremum and infimum (and hence their associated cost functions), and we also show that the algorithm for calculating the supremum must terminate after no more than T steps, where T is the number of observations. Therefore, this algorithm also provides a direct way to test for Cournot rationalizability and its role is loosely analogous to that of the generalized axiom of revealed preference (GARP), which gives a direct way to check for utility maximization in the context of consumer choice.

[2]In this respect, our test is different from those developed for general equilibrium models (like Brown and Matzkin (1996) or Kubler (2003)), which are nonlinear and computationally far more difficult.

The presence of interagent payoff effects in game-theoretic models could mean that the collection of data sets that is consistent with the model is very large, even if testable restrictions exist in principle. So it is not clear that our revealed preference test has the power to reject real data. To check that it is empirically viable, we apply (in Section 6) the test to the oil-producing countries both within and outside the Organization of Petroleum Exporting Countries (OPEC). We show that the Cournot hypothesis is clearly rejected by the data.

It is straightforward to extend the basic analysis to account for observable cost shifters (like input prices) that may affect the firms' cost functions. In Section 4, we construct a test for the Cournot hypothesis when firms face *common* shocks to marginal costs that the researcher cannot observe. The test for this model is striking in that it resembles the acyclicality tests (such as GARP) that are standard in the revealed preference analysis of consumption and production.

The investigation of market structure is of course a principal concern in empirical IO. We see our revealed preference approach as complementary to the one typically adopted in that literature. Our analysis is very parsimonious in its assumptions: we completely avoid parametric assumptions on cost and demand functions, and while the empirical IO literature typically assumes that the researcher can observe variables that are known to shift or twist the demand curves, none of that is required in our approach. However, because we make so few assumptions, we cannot do what is typically done in that literature, which is to measure the level of collusion in an industry via estimates of conduct parameters (even if we could test the Cournot hypothesis itself). The distinction between the two approaches is discussed in greater detail in Section 5.

In spite of this limitation, our revealed preference test can still be useful to antitrust authorities who consider the Cournot equilibrium as a benchmark (which is natural because firms that are not colluding ought to be playing the static Nash equilibrium at each time period). Since the test allows for great freedom in choosing the rationalizing cost and demand curves, a *rejection* of the Cournot hypothesis by this test provides very robust evidence that firm interaction is taking a more complicated (and possibly collusive) form, which could then provide a basis for further analysis through a more elaborate model. It also helps that our test requires minimal information and is very easy to implement.

Related Literature

Forges and Minelli (2009) extended Afriat's theorem to constraint sets that need not be classical budget sets and need not even be convex; the authors also pointed out that their results can be applied to games in which each player's constraint set is dependent on the actions of other players. A Cournot game belongs to this class, since the output decisions of other firms affect each firm's residual demand curve, which can be thought of as a constraint over which the firm chooses its output and the market price. However, unlike the Forges and Minelli's setup, the market demand function (and hence each player's residual demand curve) is not fully specified as part of the data in our context. Furthermore, each player's objective function is a profit function, which has a specific functional form.

The testable implications of equilibrium behavior in abstract games have been investigated by Sprumont (2000), Ray and Zhou (2001), and Galambos (2011). These papers differ from ours in two ways. First, in their work, payoff functions remain fixed and the variability in the data arises from each player being constrained to choose from different subsets of available strategies. On the other hand, our paper is most naturally understood as one in which the payoff functions are changing across observations (because of changes to demand). A second difference is that their necessary and sufficient conditions are all developed on data sets that are sufficiently complete in some sense (in this regard, Galambos's assumptions are the weakest); in our case, the data set can be as small as a single observation.[3]

2. COURNOT RATIONALIZABILITY

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. We index the observations by $t \in \mathcal{T} = \{1, 2, \ldots, T\}$. For each t, the industry price P_t and the output of each firm $(Q_{i,t})_{i\in\mathcal{I}}$ are observed, so the set of observations (or data set, for short) $\mathcal{O}$ can be written as $\{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$. We require $P_t > 0$ and $Q_{i,t} > 0$ for all (i, t), and denote the aggregate output of the industry at observation t by $Q_t = \sum_{i\in\mathcal{I}} Q_{i,t}$.

We say that the data set is Cournot rationalizable if each observation can be explained as a Cournot equilibrium that arises from a different market demand function, keeping the cost function of each firm fixed across observations, and with the demand and cost functions obeying certain regularity properties. By a *cost function* of firm i, we mean a continuous and increasing function[4] $\bar{C}_i : \mathbb{R}_+ \to \mathbb{R}$ that takes nonnegative values. The market inverse demand function $\bar{P}_t : \mathbb{R}_+ \to \mathbb{R}$ (for each t) is said to be *downward sloping* if it is differentiable at all $q > 0$, with $\bar{P}_t'(q) \le 0$; if the last inequality is strict, we say that $\bar{P}_t$ is *strictly* downward sloping. Formally, the data set $\mathcal{O} = \{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ is

[3]This difference is loosely analogous to that between the setting of Afriat's theorem and integrability results that recover utility functions from demand functions defined in an open neighborhood of prices.

[4]By "increasing" we mean here that if $q' > q$, then $\bar{C}_i(q') \ge \bar{C}_i'(q)$; in other words, we mean what is sometimes called weakly increasing or nondecreasing. Since most of the inequalities in this paper are weak, this terminology is appropriate for our purposes.

Cournot rationalizable (or *consistent with the Cournot model*) if there exist cost functions $\bar{C}_i$ for each firm i and downward-sloping demand functions $\bar{P}_t$ for each observation t such that

(i) $\bar{P}_t(Q_t) = P_t$ and

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

Condition (i) says that the inverse demand function must agree with the observed data 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. If $\mathcal{O}$ permits a rationalization where $\bar{P}_t$ are strictly downward sloping at each t, we say that $\mathcal{O}$ is *strictly* Cournot rationalizable.

Note that the question we are posing differs markedly from empirical studies of market structure in at least one respect: we are making virtually no assumptions about the behavior of demand; neither are we making any attempt to estimate demand behavior from observations on output, price, or other factors that could potentially influence demand. Put another way, we are posing the question in a way that strongly favors the Cournot model by allowing the observer to explain the data using any demand curve, as long as it passes through the observed price and output (P_t and Q_t) at each observation; in particular, no restrictions are placed on the slope of this curve, apart from requiring it to be nonpositive.

In such a setting, it should not be surprising if Cournot rationalizability alone imposes effectively no restrictions on the data set. The following result, which we prove in the Appendix, states this formally. The result is stated for a *generic* set of observations $\mathcal{O} = \{[P_t, (Q_{i,t})_{i \in \mathcal{I}}]\}_{t \in \mathcal{T}}$; by this, we mean that for all $i \in \mathcal{I}$, $Q_{i,t} \neq Q_{i,t'}$ whenever $t \neq t'$.[5]

THEOREM 1: *Suppose that $\mathcal{O} = \{[P_t, (Q_{i,t})_{i \in \mathcal{I}}]\}_{t \in \mathcal{T}}$ is a generic set of observations. Then $\mathcal{O}$ is strictly Cournot rationalizable and the firms' cost functions can be chosen to be strictly increasing and* $\mathbf{C}^2$.

Theorem 1 holds partly because, as we already pointed out, virtually no assumptions are made on the behavior of demand in the data set. But it is also crucial to this result that each firm's cost function is allowed to take on virtually any form, as long as it is increasing; if further conditions are imposed on

[5]It is possible for a nongeneric data set to be inconsistent with a Cournot model with differentiable cost functions, even if we allow for nonincreasing marginal costs. Consider the following nongeneric data set: at observation t, $P_t = 4$, $Q_{i,t} = 1$, and $Q_{j,t} = 4$, and at observation t', $P_{t'} = 2$, $Q_{i,t'} = 1$, and $Q_{j,t'} = 1$. Suppose that the firms are playing a Cournot game. Then the first-order condition of firm j at observation t tells us that $\bar{C}_j'(4) = 4\bar{P}_t'(5) + \bar{P}_t(5)$. Since marginal cost is positive and $\bar{P}_t(5) = 4$, we obtain $\bar{P}_t'(5) > -1$. Using the first-order condition for firm i, we see that $\bar{C}_i'(1) = \bar{P}_t'(5) + \bar{P}_t(5) > 3$. It is impossible for $\bar{C}_i'(1) > 3$, since firm i at observation t' is also producing 1 but the market price at t' is just 2.

the cost function, it is no longer true that "anything goes." For example, consider the case where each firm has constant marginal costs. It is well known that a firm with the larger market share in a Cournot equilibrium must have *strictly* lower marginal cost than a smaller firm, unless the two firms have identical marginal costs and the market price equals that marginal cost. This implies that, provided observed prices differ across observations, the *ranking* of firms cannot change if their marginal costs are constant, no matter how demand shifts from one observation to the next. Therefore, any data set where there is a simultaneous change in prices and in firm output rankings will not be compatible with a Cournot model with firms having constant marginal costs.

Generalizing from this observation, we could consider firms with convex cost functions, that is, with increasing marginal costs. This is a standard assumption in theoretical and econometric work because it helps to make the firm's optimization problem tractable and it is also a plausible assumption in many settings. It turns out that Cournot rationalizability with this assumption on cost functions imposes meaningful restrictions on a data set. Our main goal in this paper is to formulate and understand those restrictions.

Cournot Rationalizability With Convex Cost Functions

We begin by examining conditions on a data set that are *necessary* for it to be consistent with a Cournot model with convex cost functions. Suppose that $\mathcal{O} = \{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ is Cournot rationalizable by demand functions $\{\bar{P}_t\}_{t\in\mathcal{T}}$ and convex cost functions $\{\bar{C}_i\}_{i\in\mathcal{I}}$. We denote by $\bar{C}_i'(Q_{i,t})$ the set of subgradients of $\bar{C}_i$ at $Q_{i,t}$; the nonemptiness of this set is guaranteed by the convexity of $\bar{C}_i$. At observation t, firm i chooses q_i to maximize its profit given the output of the other firms; at its optimal choice, $Q_{i,t}$, the first-order conditions say that there is $\delta_{i,t} \geq 0$ contained in $\bar{C}_i'(Q_{i,t})$ such that

$$Q_{i,t}\bar{P}_t'(Q_t) + \bar{P}_t(Q_t) - \delta_{i,t} = 0.$$

Using condition (i) in our definition of Cournot rationalizability, it follows that the array $\{\delta_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ must obey the following condition, which we refer to as the *common ratio property*: for every $t \in \mathcal{T}$,

$$\frac{P_t - \delta_{1,t}}{Q_{1,t}} = \frac{P_t - \delta_{2,t}}{Q_{2,t}} = \cdots = \frac{P_t - \delta_{I,t}}{Q_{I,t}} \geq 0. \tag{1}$$

This holds because the first-order condition guarantees that $(P_t - \delta_{i,t})/Q_{i,t} = -\bar{P}_t'(Q_t)$, and the latter is nonnegative and independent of i. We refer to (1) as the *strict* common ratio property when the inequality is strictly positive; this holds if the $\bar{P}_t$ is strictly downward sloping. Since each firm's cost function is convex, the array $\{\delta_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ must also display increasing marginal costs, that is, for every $i \in \mathcal{I}$, $\delta_{i,t'} \geq \delta_{i,t}$ whenever $Q_{i,t'} > Q_{i,t}$. We refer to this as the *co-monotone property*. Equivalently, this property can be stated as a linear inequality $(\delta_{i,t'} - \delta_{i,t})(Q_{i,t'} - Q_{i,t}) \geq 0$.

In Example 1 below, we use the interplay of these properties to show that certain data sets are not consistent with the Cournot model with convex cost functions. The example also illustrates the potential usefulness of our test for a market regulator: the two firms exhibit behavior that seems collusive—between t and t' they simultaneously reduce outputs, leading to an increase in the market price—but before a regulator investigates this possibility, she may wish to eliminate (with a minimum of ancillary assumptions) the alternative that the firms are simply playing a Cournot game.

EXAMPLE 1: Consider the following observations of two firms i and j:

(i) At observation t, $P_t = 4$, $Q_{i,t} = 60$, and $Q_{j,t} = 110$.

(ii) At observation t', $P_{t'} = 10$, $Q_{i,t'} = 50$, and $Q_{j,t'} = 100$.

We claim that these observations are not Cournot rationalizable with convex cost functions. Indeed, if they are, then there is $\delta_{i,t'} \in \bar{C}'_i(Q_{i,t'})$ and $\delta_{j,t'} \in \bar{C}'_j(Q_{j,t'})$ such that

$$\delta_{i,t'} = P_{t'} - (P_{t'} - \delta_{j,t'})\frac{Q_{i,t'}}{Q_{j,t'}} \geq P_{t'}\left(1 - \frac{Q_{i,t'}}{Q_{j,t'}}\right), \tag{2}$$

where the equality follows from the common ratio property and the inequality follows from the assumption that marginal cost is positive. Substituting in the numbers given, we obtain $\delta_{i,t'} \geq 5$, where $\delta_{i,t'} \in \bar{C}'_i(50)$. Since firm i has increasing marginal costs, the co-monotone property holds, so $\delta_{i,t} \geq 5$ since $Q_{i,t} > Q_{i,t'}$. However, this means that $\delta_{i,t}$ exceeds $P_t = 4$, which violates the common ratio property (at t).

Not only is Cournot rationalizability with convex cost functions refutable, the hypothesis of strict Cournot rationalizability with convex cost functions can be refuted even when price information is not available. (A realistic environment where that occurs is described in Section 4.) We show this in the next example, which exploits the following important consequence of the strict common ratio property: the firm with strictly lower output must have strictly lower marginal cost.

EXAMPLE 2: Suppose that at observation t, firm i produces 20 and firm j produces 15. At another observation t', firm i produces 15 and firm j produces 16. Suppose, contrary to our claim, that these observations are strictly Cournot rationalizable with convex cost functions. In that case, by the strict common ratio property, observation t tells us that there are $\delta_{i,t} \in \bar{C}'_i(20)$ and $\delta_{j,t} \in \bar{C}'_j(15)$ such that $\delta_{i,t} < \delta_{j,t}$. At observation t', firm i produces 15, which is less than its output at t, so by the co-monotone property, $\delta_{i,t'} \leq \delta_{i,t}$; similarly, the co-monotone property guarantees that $\delta_{j,t'} \geq \delta_{j,t}$. Putting these together, we obtain

$$\delta_{i,t'} \leq \delta_{i,t} < \delta_{j,t} \leq \delta_{j,t'},$$

but this violates the strict common ratio property since it means that at observation t', firm j has larger output *and* higher marginal cost compared to firm j at t'.

Notice that the data do not violate Cournot rationalizability. Given that the market prices are not provided, we are free to choose its values, so let us suppose that the market price at both t and t' is $m > 0$. Then the data can be explained as a Cournot outcome in the case where both firms have the same constant marginal cost of m and the demand curve at both t and t' is flat, with $\bar{P}_t(q) = \bar{P}_{t'}(q) = m$ for all $q > 0$.[6]

In both Examples 1 and 2, the common ratio property is first used to extract information about marginal costs at a firm's observed output and the co-monotone property is then applied to extend that information from one output level to a whole range of output levels. Comparing these examples with Theorem 1, notice that in the latter setting, the co-monotone property does not hold because cost functions need not be convex; while it is still possible to use the common ratio property at each observation to derive information on marginal costs across firms, this information stays at the observed output levels and cannot be extended in a way that leads to mutually inconsistent observations. So instead we obtain Theorem 1, which says that anything goes if cost functions are not restricted.

The next theorem is the main result of this section and shows that a set of observations is Cournot rationalizable with convex cost functions if and only if there exist nonnegative numbers $\{\delta_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ that obey the common ratio and co-monotone properties.[7] These two properties both impose linear conditions, so the existence (or otherwise) of $\{\delta_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ is a computationally straightforward linear program.

THEOREM 2: *The following statements on* $\{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ *are equivalent*:

(A) *The set of observations is Cournot rationalizable with convex cost functions*.

(B) *There exists a set of nonnegative numbers* $\{\delta_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ *that satisfy the common ratio property and the co-monotone property*.

[6]Essentially the same argument guarantees that *any* data set with output information but no price information is Cournot rationalizable. For the data set in this example, the reader can easily check that Cournot rationalizability with convex cost functions will require the market prices at t and t' to be identical. In other words, were price information to be available and $P_t \neq P_{t'}$, then the data set is not Cournot rationalizable with convex cost functions.

[7]Theorem 5 in the next section provides an equivalent test that is free of unknowns.

REMARK 1: When a data set is Cournot rationalizable, the convex cost functions $\bar{C}_i$ (for all $i \in \mathcal{I}$) that rationalize the data can always be chosen to obey $\bar{C}_i(0) = 0$.

REMARK 2: A straightforward variation of Theorem 2 will say that a data set is *strictly* Cournot rationalizable with convex cost functions if and only if there exist nonnegative numbers $\{\delta_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ that satisfy the strict common ratio property and the co-monotone property. Furthermore, in this case, we can guarantee that at every observation t, $[P_t, (Q_{i,t})_{i\in\mathcal{I}}]$ is the *unique* Cournot equilibrium.[8]

It is worth pointing out that Theorem 2 is useful even in situations where the output of one or more firms is missing from the data set. This is because if all of the firms in an industry are playing a Cournot game, then any subset of firms whose outputs *are* observed must also be playing a Cournot game (with the residual demand function as their "market" demand function), and the latter hypothesis can be tested using the theorem.[9]

Our proof of Theorem 2 uses two lemmas; the first one provides an explicit construction of the demand curve needed to rationalize the data at any observation t, while the second lemma provides a way to construct a cost curve for each firm that obeys stipulated conditions on marginal cost.

LEMMA 1: *Suppose that, at some observation t, there are positive scalars $\{\delta_{i,t}\}_{i\in\mathcal{I}}$ such that equation (1) is satisfied and that there are convex cost functions $\bar{C}_i$ with $\delta_{i,t} \in \bar{C}_i'(Q_{i,t})$. Then there exists a downward-sloping demand function $\bar{P}_t$ such that $\bar{P}_t(Q_t) = P_t$ and, with each firm i having the cost function $\bar{C}_i$, $\{Q_{i,t}\}_{i\in\mathcal{I}}$ constitutes a Cournot equilibrium.*

Furthermore, if (1) is satisfied with a strict inequality, then $\bar{P}_t$ can be chosen to be strictly downward sloping and $\{Q_{i,t}\}_{i\in\mathcal{I}}$ constitutes the unique *Cournot equilibrium.*

PROOF: We define $\bar{P}_t$ by $\bar{P}_t(Q) = a_t - b_t Q$, where $b_t = [P_t - \delta_{i,t}]/Q_{i,t}$—notice that this is well defined because of (1)—and we choose a_t such that $\bar{P}_t(Q_t) = P_t$. Firm i's decision is to choose $q_i \geq 0$ to maximize $\Pi_{i,t}(q_i) = q_i\bar{P}_t(q_i + \sum_{j\neq i} Q_{j,t}) - \bar{C}_i(q_i)$. This function is concave, so an output level is

[8]Example 2 gives a data set that is Cournot rationalizable (if $P_t = P_{t'} = m > 0$) but not strongly Cournot rationalizable. Note that the data *cannot* be rationalized as unique Cournot equilibria, since the rationalizing demand curve and the marginal cost curve must both be flat and equal to m, and so any output combination constitutes an equilibrium.

[9]In this regard, it is quite different from the inequality conditions of Afriat's theorem, which, unless preferences are separable, become vacuous when there is missing data (see Varian (1988)).

optimal if and only if it obeys the first-order condition. Since $\delta_{i,t} \in \bar{C}_i'(Q_{i,t})$ and since $\bar{P}_t'(Q_t) = -b_t$, a supergradient[10] of $\Pi_{i,t}$ at $Q_{i,t}$ is

$$Q_{i,t}\bar{P}_t'(Q_t) + \bar{P}_t(Q_t) - \delta_{i,t} = -Q_{i,t}\frac{[P_t - \delta_{i,t}]}{Q_{i,t}} + P_t - \delta_{i,t} = 0.$$

So we have shown that $Q_{i,t}$ is profit-maximizing for firm i at observation t.

If (1) holds with strict inequality, $b_t = [P_t - \delta_{i,t}]/Q_{i,t} > 0$ and the demand curve $\bar{P}_t$ is strictly downward sloping. Suppose that there is another equilibrium $\{Q_i'\}_{i\in\mathcal{I}}$ besides $\{Q_{i,t}\}_{i\in\mathcal{I}}$ and suppose that aggregate output is not lower than $Q_t = \sum_{i\in\mathcal{I}} Q_{i,t}$ (the case where aggregate output is lower has a similar proof). Hence, the market price at this equilibrium, P', must be no higher than P_t and there is some firm k such that $Q_k' > Q_{k,t}$. Since firm k has increasing marginal cost, we obtain

$$(3) \qquad 0 < b_t = \frac{(P' - \bar{C}_k'(Q_k'))}{Q_k'} < \frac{(P_t - \bar{C}_k'(Q_{k,t}))}{Q_{k,t}} = b_t,$$

where the equalities follow from the first-order condition. Clearly, (3) is impossible.[11] *Q.E.D.*

LEMMA 2: *Suppose that for some firm i, there are positive scalars $\{\delta_{i,t}\}_{t\in\mathcal{T}}$ that are increasing with $Q_{i,t}$. Then there exists a convex cost function $\bar{C}_i$ such that $\delta_{i,t} \in \bar{C}_i'(Q_{i,t})$.*

PROOF: Define $\hat{Q} = \{q_i \in \mathbb{R}_+ : q_i = Q_{i,t}$ for some observation $t\}$; $\hat{Q}$ consists of those output levels actually chosen by firm i at some observation. Since $\{\delta_{i,t}\}_{t\in\mathcal{T}}$ are increasing with $Q_{i,t}$, it is possible to construct a nonnegative and increasing function $\bar{m}_i : \mathbb{R}_+ \to \mathbb{R}$ with the following properties: (a) for any output $\hat{q} \in \hat{Q}$, set $\bar{m}_i(\hat{q}) = \max\{\delta_{i,t} : Q_{i,t} = \hat{q}\}$; (b) for any $\hat{q} \in \hat{Q}$, $\lim_{q\to\hat{q}^-} \bar{m}_i(q) = \min\{\delta_{i,t} : Q_{i,t} = \hat{q}\}$; (c) $\bar{m}_i$ is continuous at all $q \notin \hat{Q}$. The function $\bar{m}_i$ is piecewise continuous with a discontinuity at $\hat{q} \in \hat{Q}$ if and only if the set $\{\delta_{i,t} : Q_{i,t} = \hat{q}\}$ is nonsingleton. Define $\bar{C}_i : \mathbb{R} \to \mathbb{R}$ by

$$(4) \qquad \bar{C}_i(q) = \int_0^q \bar{m}_i(s)\,ds.$$

This function is increasing because $\bar{m}_i$ is nonnegative and it is convex because $\bar{m}_i$ is increasing. Last, (a) and (b) guarantee that $\delta_{i,t} \in \bar{C}_i'(Q_{i,t})$. *Q.E.D.*

[10]By a *supergradient* of a concave function F at a point, we mean the subgradient of the convex function $-F$ at the same point.

[11]More generally, it is known that the Cournot equilibrium is unique if firms have convex cost functions and demand functions are concave (see Szidarovsky and Yakowitz (1977)).

PROOF OF THEOREM 2 (AND REMARKS): To see that (A) implies (B), suppose that the data are rationalized with demand functions $\{\bar{P}_t\}_{t\in\mathcal{T}}$ and cost functions $\{\bar{C}_i\}_{i\in\mathcal{I}}$. We have already shown that the first-order condition guarantees the existence of $\delta_{i,t} \in \bar{C}_i'(Q_{i,t})$ obeying the common ratio property (1). Since $\bar{C}_i$ is convex, $\{\delta_{i,t}\}_{t\in\mathcal{T}}$ is increasing with $Q_{i,t}$.

The fact that (B) implies (A) is an immediate consequence of Lemmas 1 and 2. Note that the rationalizing cost function $\bar{C}_i$, as defined by (4), obeys $\bar{C}_i(0) = 0$, which confirms Remark 1. Furthermore, in the case where the strict common ratio property holds, Lemma 1 also guarantees that, for all t, $\{Q_{i,t}\}_{i\in\mathcal{I}}$ constitutes the *unique* Cournot equilibrium. This establishes the claim in Remark 2. *Q.E.D.*

Sometimes it is convenient to consider rationalizations where each firm's cost functions are differentiable (so kinks on the cost curves are not allowed). This can be characterized by strengthening the condition imposed on $\{\delta_{i,t}\}_{t\in\mathcal{T}}$ in Theorem 2; for any firm i, we say that $\{\delta_{i,t}\}_{t\in\mathcal{T}}$ has the *fine co-monotone property* if $\delta_{i,t'} \geq \delta_{i,t}$ whenever $Q_{i,t} \geq Q_{i,t'}$ and $\delta_{i,t'} = \delta_{i,t}$ whenever $Q_{i,t} = Q_{i,t'}$. The next result is proved in the Appendix.

THEOREM 3: *The following statements on $\mathcal{O} = \{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ are equivalent*:

(A) *The set $\mathcal{O}$ is Cournot rationalizable with C^2 convex cost functions for all firms*.[12]

(B) *There exists a set of nonnegative numbers $\{\delta_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ that satisfy the common ratio property and the fine co-monotone property*.

3. SET IDENTIFICATION OF MARGINAL COSTS

Beginning with Rosse (1970), there is a large literature in empirical industrial organization that seeks to recover marginal costs of firms from observed prices and quantities. Unlike the framework of this paper, it is typical in this literature to make parametric assumptions for both market demand and firms' cost functions. In this section, we design a simple procedure by which the set of marginal costs that rationalizes the data can be identified. Given that only convex cost functions will be considered from this point on, we simply refer to such data sets as *Cournot rationalizable* or even *rationalizable*; the restriction to convex cost functions is implicit.

For a Cournot rationalizable data set $\mathcal{O} = \{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$, let $\{\delta_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ be a solution to the linear program in statement (B) of Theorem 2. We may

[12]It is clear from the proof of this result that the cost functions could in fact be chosen to be differentiable to any degree, but there is no particular need to go beyond C^2, which is sufficient to ensure the differentiability of the marginal cost function.

represent this solution as a vector δ in R_+^{IT} and represent the set of solution vectors by $M \subset R_+^{IT}$. It is clear that M *set-identifies* the family of convex cost functions that are compatible with Cournot rationalizability. By this we mean that $\mathcal{O}$ is rationalizable with convex cost functions $\{\bar{C}_i\}_{i\in\mathcal{I}}$ if and only if there is $\delta_{i,t} \in \bar{C}_i'(Q_{i,t})$ such that $\delta = (\delta_{i,t})_{(i,t)\in\mathcal{I}\times\mathcal{T}} \in M$. This follows immediately from Lemmas 1 and 2. In other words, Theorem 2 does not simply provide a test of Cournot rationalizability: solving the linear program in statement (B) of the theorem also allows us to recover all the cost information on firms that could be recovered from the model. For this reason, it is important that we have a good understanding of the set M; the following result tells us more about this set.

THEOREM 4: *Suppose that $\mathcal{O} = \{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ is Cournot rationalizable with convex cost functions. Then the solution set M is a nonempty compact and convex sublattice of R_+^{IT}. Consequently, $\delta^U = \sup M$ and $\delta^L = \inf M$ exist and are contained in M.*[13]

PROOF: The existence of δ^U, $\delta^L \in M$ is guaranteed as long as M is a compact sublattice (see Topkis (1998, Theorem 2.3.1)). Since M is defined as a set that obeys a family of weak linear inequalities, it must be closed and convex. It is also bounded since $0 \leq \delta_{i,t} \leq P_t$. Therefore M is compact. The fact that M is a sublattice of R_+^{IT} follows from direct verification. Indeed, suppose δ and $\tilde{\delta}$ are both in M; we show that $\delta \vee \tilde{\delta}$ and $\delta \wedge \tilde{\delta}$ are in M by checking that they obey the co-monotone and common ratio properties. Suppose that for some firm k, $Q_{k,t} > Q_{k,t'}$; the co-monotone property implies that $\delta_{k,t} \geq \delta_{k,t'}$ and $\tilde{\delta}_{k,t} \geq \tilde{\delta}_{k,t'}$. Therefore, $\max\{\delta_{k,t}, \tilde{\delta}_{k,t}\} \geq \delta_{k,t'}$ and $\max\{\delta_{k,t}, \tilde{\delta}_{k,t}\} \geq \tilde{\delta}_{k,t'}$. Consequently, $\max\{\delta_{k,t}, \tilde{\delta}_{k,t}\} \geq \max\{\delta_{k,t'}, \tilde{\delta}_{k,t'}\}$. A similar argument will guarantee that $\min\{\delta_{k,t}, \tilde{\delta}_{k,t}\} \geq \min\{\delta_{k,t'}, \tilde{\delta}_{k,t'}\}$. We conclude that $\delta \vee \tilde{\delta}$ and $\delta \wedge \tilde{\delta}$ obey the co-monotone property.

Now suppose that, for fixed t, $\{\delta_{i,t}\}_{i\in\mathcal{I}}$ and $\{\tilde{\delta}_{i,t}\}_{i\in\mathcal{I}}$ both obey the common ratio property. With no loss of generality, suppose that for some firm j, $\delta_{j,t} \geq \tilde{\delta}_{j,t}$. The common ratio property guarantees that $\delta_{i,t} \geq \tilde{\delta}_{i,t}$ for *all* $i \in \mathcal{I}$. Therefore, $\{\max\{\delta_{i,t}, \tilde{\delta}_{i,t}\}\}_{i\in\mathcal{I}}$ obeys the common ratio property since it is identical to $\{\delta_{i,t}\}_{i\in\mathcal{I}}$ and $\{\min\{\delta_{i,t}, \tilde{\delta}_{i,t}\}\}_{i\in\mathcal{I}}$ obeys the common ratio property since it coincides with $\{\tilde{\delta}_{i,t}\}_{i\in\mathcal{I}}$. *Q.E.D.*

Theorem 4 tells us that M contains its supremum (δ^U) and infimum (δ^L), which represent the highest and lowest marginal costs (at each firm's observed

[13]Suppose the data set is *strictly* Cournot rationalizable and let M^* be the corresponding set of solutions (see Remark 2 following Theorem 2). It is straightforward to check that M^* is also a convex lattice, though it is not necessarily closed. The closure of M^* equals M because any strict convex combination of a point in M^* and a point in M must be in M^*. It follows that $\sup M^* = \delta^U$ and $\inf M^* = \delta^L$.

outputs) that are consistent with a Cournot rationalization of the data. For any Cournot rationalizable data set, there are intuitive algorithms for finding δ^U and δ^L. The following *upper bound algorithm* (or UB algorithm) constructs, at each step, *a decreasing sequence of upper bounds on M that converges to δ^U*.

UPPER BOUND ALGORITHM:
1. Initialize $\delta^{\text{ub}}_{i,t} = P_t$ for all (i, t). Go to Step 3.
2. Set $\delta^{\text{ub}}_{i,t} = \gamma^{\text{ub}}_{i,t}$ for all (i, t).
3. For each (i, t), set $\gamma^{\text{ub}}_{i,t} = \min\{\min_{\{t' \neq t: Q_{i,t'} > Q_{i,t}\}}\{\delta^{\text{ub}}_{i,t'}\}, \delta^{\text{ub}}_{i,t}\}$.[14]
4. Let $\Lambda_t = \max_j\{\frac{P_t - \gamma_{j,t}}{Q_{j,t}}\}$. For each (i, t), set $\gamma^{\text{ub}}_{i,t} = P_t - \Lambda_t Q_{i,t}$.
5. If any $\gamma^{\text{ub}}_{i,t} < 0$, then stop and output that the data set is not rationalizable. If $\delta^{\text{ub}}_{i,t} = \gamma^{\text{ub}}_{i,t}$ for all (i, t), then stop and output the bounds. Else, go to Step 2.

This algorithm is very easy to understand. In Step 1, we set $\delta^{\text{ub}}_{i,t} = P_t$; clearly, this gives the highest possible marginal cost of firm i at $Q_{i,t}$. The vector $\delta^{\text{ub}} = (\delta^{\text{ub}}_{i,t})_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ is an upper bound of M and it obeys the common ratio property. However, the co-monotone property need not hold. In Step 3, we use the fact that firm i has increasing marginal costs (in other words, the co-monotone property) to adjust the upper bound downward to $\gamma^{\text{ub}}_{i,t}$. Note that γ^{ub} is again an upper bound of M. At the end of Step 3 the co-monotone property holds but the common ratio property need not be satisfied. In Step 4, the upper bound is adjusted downward again, this time using the common ratio property. This procedure is repeated in the second iteration, where the upper bound is first lowered using the co-monotone property and then lowered again using the common ratio property. Given a Cournot rationalizable data set, the vector δ^{ub}_n, obtained after Step 2 of the nth iteration, is an upper bound of M and is decreasing in n. It must converge to some limit, and at the limit, both the common ratio and the co-monotone properties will hold. Therefore, the limit is itself an element of M and we obtain

$$\lim_{n\to\infty} \delta^{\text{ub}}_n = \delta^U. \tag{5}$$

The following example shows the algorithm at work.

EXAMPLE 3: Consider the following observations of two firms i and j:
(i) At observation t_1, $P_{t_1} = 28$, $Q_{i,t_1} = 1$, and $Q_{j,t_1} = 1$.
(ii) At observation t_2, $P_{t_2} = 26$, $Q_{i,t_2} = 5$, and $Q_{j,t_2} = 10$.
(iii) At observation t_3, $P_{t_3} = 24$, $Q_{i,t_3} = 10$, and $Q_{j,t_3} = 20$.
(iv) At observation t_4, $P_{t_4} = 22$, $Q_{i,t_4} = 15$, and $Q_{j,t_4} = 90$.
(v) At observation t_5, $P_{t_5} = 20$, $Q_{i,t_5} = 20$, and $Q_{j,t_5} = 100$.

[14]Set $\min_{\{t' \neq t: Q_{i,t'} > Q_{i,t}\}}\{\delta^{\text{ub}}_{i,t'}\} = \infty$ if $\{t' \neq t : Q_{i,t'} > Q_{i,t}\}$ is an empty set.

Below, we display the upper bounds obtained in Steps 1 and 3 of the first iteration of the UB algorithm. The entry in row t and column k of the matrix corresponds to the upper bound on the marginal cost of firm k at output $Q_{k,t}$:

$$\delta_1^{\text{ub}} = \begin{bmatrix} 28 & 28 \\ 26 & 26 \\ 24 & 24 \\ 22 & 22 \\ 20 & 20 \end{bmatrix} \to \gamma_1^{\text{ub}} = \begin{bmatrix} 20 & 20 \\ 20 & 20 \\ 20 & 20 \\ 20 & 20 \\ 20 & 20 \end{bmatrix}.$$

The entries in γ_1^{ub} are all equal to 20 because the output of both firms is greatest at t_5, so the price at t_5 must bound the marginal cost at all observed output levels. After Step 3, the co-monotone property holds, but clearly not the common ratio property; applying that property in Step 4 lowers the marginal cost of (the smaller) firm 2 further and gives us δ_2^{ub}. Repeating this process, we obtain δ^U after five iterations:

$$\delta_2^{\text{ub}} = \begin{bmatrix} 20 & 20 \\ 20 & 14 \\ 20 & 16 \\ 20 & 10 \\ 20 & 20 \end{bmatrix} \to \delta_3^{\text{ub}} = \begin{bmatrix} 10 & 10 \\ 18 & 10 \\ 17 & 10 \\ 20 & 10 \\ 20 & 20 \end{bmatrix} \to \delta_4^{\text{ub}} = \begin{bmatrix} 10 & 10 \\ 17 & 8 \\ 17 & 10 \\ 20 & 10 \\ 20 & 20 \end{bmatrix}$$

$$\to \delta_5^{\text{ub}} = \delta^U = \begin{bmatrix} 8 & 8 \\ 17 & 8 \\ 17 & 10 \\ 20 & 10 \\ 20 & 20 \end{bmatrix}.$$

There is a similar *lower bound* (LB) *algorithm* that gives, at each step, *an increasing sequence of lower bounds on M that converges to δ^L.*

LOWER BOUND ALGORITHM:

1. Initialize $\delta_{i,t}^{\text{lb}} = \gamma_{i,t}^{\text{lb}} = 0$ for all (i, t). Go to Step 3.
2. Set $\delta_{i,t}^{\text{lb}} = \gamma_{i,t}^{\text{lb}}$ for all (i, t).
3. For each (i, t), set $\gamma_{i,t}^{\text{lb}} = \max\{\max_{\{t'\neq t:Q_{i,t'}<Q_{i,t}\}}\{\delta_{i,t'}^{\text{lb}}\}, \delta_{i,t}^{\text{lb}}\}$.[15]
4. Let $\Lambda = \min_j\{\frac{P_t-\gamma_{j,t}}{Q_{j,t}}\}$. For each i, t, set $\gamma_{i,t}^{\text{lb}} = P_t - \Lambda Q_{i,t}$.
5. If any $\gamma_{i,t}^{\text{lb}} > P_t$, then stop and output that the data set is not rationalizable. If $\delta_{i,t}^{\text{lb}} = \gamma_{i,t}^{\text{lb}}$ for all (i, t), then stop and output the bounds. Else, go to Step 2.

In Step 1, we set $\delta_{i,t}^{\text{lb}} = 0$, which is trivially a lower bound on all marginal costs. This bound remains unchanged in Step 3. In Step 4, the algorithm uses

[15]Set $\max_{\{t'\neq t:\,Q_{i,t'}<Q_{i,t}\}}\{\delta_{i,t'}^{\text{lb}}\} = 0$ if $\{t' \neq t : Q_{i,t'} < Q_{i,t}\}$ is an empty set.

the common ratio property to raise the lower bounds on marginal cost for each firm, using the fact that the marginal cost of the firm with the largest output at each observation cannot fall below zero. The vector $\gamma^{\text{lb}} = (\gamma^{\text{lb}}_{i,t})_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ obtained after Step 4 is a lower bound on M that obeys the common ratio property but not necessarily the co-monotone property. In Step 3 of the second iteration, the lower bound is adjusted upward again, this time using the co-monotone property; this is followed by another upward adjustment using the common ratio property, and so forth. Given a Cournot rationalizable data set, the vector δ^{lb}_n, obtained after Step 2 of the nth iteration, is a lower bound of M, and is an increasing and bounded sequence (in n). It must converge to some limit, and at the limit, both the common ratio and the co-monotone properties will hold. Therefore, the limit is itself an element of M and we obtain

$$\lim_{n\to\infty} \delta^{\text{lb}}_n = \delta^L. \tag{6}$$

While the UB and LB algorithms allow us to find δ^U and δ^L in a very intuitive way, our discussion so far does not guarantee that the processes terminate after some known number of iterations, so it is not clear that they could be used for testing. The next result shows that *the UB algorithm will terminate affirmatively at or before T iterations if and only if the data set is Cournot rationalizable*. (Recall that T is the number of observations in the data set.) By *terminate affirmatively*, we mean that at Step 5 of the UB algorithm, we obtain $\delta^{\text{ub}}_{i,t} = \gamma^{\text{ub}}_{i,t}$ for all i, t. Thus the UB provides an alternative way to check for Cournot rationalizability: the researcher will know whether the data set is Cournot rationalizable after no more than T iterations. In contrast to Theorem 2, the UB algorithm gives a procedure that does not require verification of the existence of real numbers that satisfy a set of inequalities. In this way, the equivalence of the UB algorithm and the test in Theorem 2 for Cournot rationalizability is analogous to the equivalence of the generalized axiom of revealed preference (GARP) and the Afriat inequalities test for utility maximization. The proof of Theorem 5 is in the Appendix. Note that Example 3 shows that the bound of T iterations is tight.

THEOREM 5: *The data set $\mathcal{O} = \{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ is Cournot rationalizable with convex cost functions if and only if the UB algorithm terminates affirmatively. Furthermore, affirmative termination must take place at or before T iterations and $\delta^H = (\tilde{\delta}^{\text{ub}}_{i,t})_{(i,t)\in\mathcal{I}\times\mathcal{T}}$, where $\tilde{\delta}^{\text{ub}}_{i,t}$ is the final value of $\delta^{\text{ub}}_{i,t}$ in the algorithm.*

REMARK 3: If $\mathcal{O}$ is not Cournot rationalizable with convex cost functions, δ^{ub} *must* stray outside R^{IT}_+ after a finite number of iterations (which can be less than or greater than T); at that point, the algorithm will terminate, but not affirmatively (see Step 5 of the UB algorithm). This is because δ^{ub} is a decreasing sequence by construction and if it always stays nonnegative (so the algorithm goes on forever), it will have a limit that must be an element of M (see (5)).

REMARK 4: It follows from Theorems 2 and 5 that the UB algorithm terminates affirmatively if and only if there is a solution to the linear program in statement (B) of Theorem 2.

4. STRICT COURNOT RATIONALIZABILITY WITH COMMON COST SHOCKS

In this section, we develop a test for Cournot rationalizability in the case where cost functions are also changing, *in addition* to changes in demand. It is quite clear that if we allow each firm's cost function to change arbitrarily across observations, then there would effectively be *no* observable restrictions on the data set. But there are at least two ways to introduce cost changes that will still lead to meaningful tests of Cournot rationalizability. One possibility is to assume that the researcher can observe the variation of the parameters (such as input prices) that have an impact on firms' marginal costs in the data set. Theorem 3 can be adapted to allow for the cost functions to change when those observed parameters change, by replacing the co-monotone property with a requirement that marginal costs are greater only when both output and the observed parameters are higher.[16] A detailed formulation of this result can be found in Carvajal, Deb, Fenske, and Quah (2010).

The other possibility, which we consider here, is to assume that cost changes are completely unobserved but *common* across the industry. Assuming that input price changes are the main source of changes to the cost functions, then common cost shocks will arise if the input price changes that occur in the data arise primarily from those inputs where there is no substantial difference in usage across firms.[17] In this case, we obtain a simple and quite striking revealed preference test on the data set.

We say that a data set $\mathcal{O} = \{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ is *strictly Cournot rationalizable with convex cost functions under common perturbations* if there are permanent convex cost functions $\bar{C}_i$ for each firm i, and, for each observation t, there are strictly downward-sloping inverse demand functions $\bar{P}_t$ passing through (Q_t, P_t) and convex cost perturbation functions f_t, such that the following statement holds: for all (i, t), $Q_{i,t}$ is firm i's best response to $\sum_{j\neq i} Q_{j,t}$ when $\bar{P}_t$ is the inverse demand function and firm i has the cost function

$$\bar{C}_{i,t}(q_i) = \bar{C}_i(q_i) + f_t(q_i). \tag{7}$$

[16]For example, the observed parameter could be the price of some major production input; in this case, marginal cost will increase in the input price if the demand for this input (as a function of the output level) is normal. If the observed parameter is a vector (for example, the prices of different inputs), then these observed parameters may not always be comparable, in which case the marginal cost functions (of each firm) across observations would not be completely ordered.

[17]To be specific, suppose that K is the set of inputs required for production and that firms' production functions are of the Leontief type, where $a_i^k(q)$ is the level of input $k \in K$ needed to produce output q by firm i. Let K' be the set of inputs whose prices do vary in the data set, while the input prices in $K \setminus K'$ are stable. Then common cost shocks require $a_i^k(q) = a_j^k(q)$ for any two firms i and j and $k \in K'$, while $a_i^k(q)$ may differ across firms for $k \in K \setminus K'$.

It is not immediately clear that this notion of rationalizability imposes meaningful restrictions on the data set, though in fact it does: it imposes an intuitive no-cycling condition on observed firms' outputs.

To see why, we first define the relations $\succeq_0$ and $\succ_0$ on $(i,t)\in\mathcal{I}\times\mathcal{T}$: for any $i\neq j$, $t\neq t'$, we say that $(i,t)\succeq_0(i,t')$ if $Q_{i,t}>Q_{i,t'}$, $(i,t)\succeq_0(j,t)$ if $Q_{j,t}\geq Q_{i,t}$, and $(i,t)\succ_0(j,t)$ if $Q_{j,t}>Q_{i,t}$. We denote the transitive closure of $\succeq_0$ by $\succeq$, that is, $(i,t)\succeq(j,t')$ if there is some sequence $(i_1,t_1),(i_2,t_2),\dots,(i_n,t_n)$ such that

$$(i,t)\succeq_0(i_1,t_1)\succeq_0(i_2,t_2)\succeq_0\cdots\succeq_0(i_n,t_n)\succeq_0(j,t').$$

Last, define $\succ$ by saying that $(i,t)\succ(j,t')$ whenever any of the relations in the sequence above could be replaced with $\succ_0$. We say that the data set $\mathcal{O}$ satisfies *marginal cost consistency* (MCC for short) if

$$(i,t)\succeq(j,t')\quad\Longrightarrow\quad(j,t')\not\succ(i,t)\quad\text{for all } i,j,t,t'.$$

For an instance of a data set that does not obey MCC, consider Example 2 in Section 2. In that example, firm i reduces output between t and t', while firm j increases output; furthermore, firm i is the larger firm at observation t, but it is the smaller firm at t'. In other words, $Q_{i,t}>Q_{i,t'}$, $Q_{i,t'}<Q_{j,t'}$, $Q_{j,t'}>Q_{j,t}$, and $Q_{j,t}<Q_{i,t}$. This implies, respectively, that

$$\begin{aligned}&(i,t)\succeq_0(i,t'),\quad (i,t')\succ_0(j,t'),\\&(j,t')\succeq_0(j,t),\quad\text{and}\quad(j,t)\succ_0(i,t),\end{aligned}$$

which is clearly inconsistent.

We now explain why marginal cost consistency must be true of any data set that is strictly Cournot rationalizable with convex cost functions under common perturbations; as we do so, it will also become clear why the term "marginal cost consistency" makes sense. Since firm i at observation t is maximizing profit at $Q_{i,t}$, the first-order condition says that

$$P_t+Q_{i,t}\bar{P}_t'(Q_t)=\delta_{i,t}+\phi_{i,t},$$

where $\delta_{i,t}\in\bar{C}_i'(Q_{i,t})$ and $\phi_{i,t}\in f_t'(Q_{i,t})$. If $(i,t)\succeq_0(i,t')$, then $Q_{i,t}>Q_{i,t'}$ and so, by the convexity of $\bar{C}_i$, $\delta_{i,t}\geq\delta_{i,t'}$. Now suppose $(i,t)\succeq_0(\succ_0)(j,t)$ so $Q_{j,t}\geq(>)$ $Q_{i,t}$; given that firm j is producing more (strictly more) than i at observation t, it must have lower (strictly lower) marginal cost, that is,

$$\delta_{i,t}+\phi_{i,t}\geq(>)\ \delta_{j,t}+\phi_{j,t}.$$

(Note that the "strict" part of this claim relies on the fact that the data are strictly Cournot rationalizable, rather than just Cournot rationalizable.) Since

f_t is convex and $Q_{i,t} \le Q_{j,t}$, we obtain $\phi_{i,t} \le \phi_{j,t}$; therefore, $\delta_{i,t} \ge (>)\ \delta_{j,t}$. To summarize, we have shown that

$$\text{if} \quad (i,t) \succeq (\succ)\ (j,t'), \quad \text{then} \quad \delta_{i,t} \ge (>)\ \delta_{j,t'}.$$

It is clear that this property implies MCC.

The next result says that MCC is both necessary and sufficient for Cournot rationalizability when common cost perturbations are allowed.

THEOREM 6: *The data set* $\mathcal{O} = \{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$ *is Cournot rationalizable with convex cost functions under common perturbations if and only if it obeys MCC.*

The proof of the sufficiency of MCC can be found in the Appendix; the proof works by showing that MCC implies that there is a solution to a particular linear program and then showing that the existence of a solution to the linear program implies rationalizability. Note that the MCC condition involves only firm outputs and not prices, so the price information in the data set is not needed for testing the hypothesis.

5. TESTING FOR COLLUSION

A major concern in empirical IO is the detection of collusive behavior. This question is related to, but distinct from, the principal focus of our paper, which is to develop revealed preference tests for Cournot behavior. This section is devoted to explaining this distinction.

Consider once again the basic environment of Sections 2 and 3, in which we developed a test for the Cournot rationalizability of a data set $\mathcal{O} = \{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$. What conditions are needed for $\mathcal{O}$ be consistent with *perfect collusion*, in the sense of all firms acting in concert to maximize joint profit? This question has a short and simple answer: *any* data set is consistent with perfect collusion. To see this, suppose that every firm has the same cost function $\bar{C}(q) = \epsilon q$. Then every output allocation is cost efficient and if firms are colluding, they will act like a monopoly with the same cost function $\bar{C}$. Choose ϵ sufficiently small so that $P_t > \epsilon$ for all $t \in \mathcal{T}$. Clearly there is a linear and downward-sloping inverse demand function $\bar{P}_t$ such that $\bar{P}_t(Q_t) = P_t$ and the monopolist's profit is maximized at Q_t.

To reinforce the message that collusion cannot be excluded in this setting, we embed the Cournot model within a model of conjectural variations, which is commonly used in empirical estimates of market power (see Bresnahan (1989)). To each firm i, we associate a scalar $\theta_i \ge 0$ that is the firm's *conduct parameter*; the output vector $(Q_i^*)_{i\in\mathcal{I}}$ constitutes a $\theta = \{\theta_i\}_{i\in\mathcal{I}}$ *conjectural variations equilibrium* (or θ-CV equilibrium for short) if

$$Q_i^* \in \arg\max_{q_i \ge 0} \left\{ q_i \bar{P}\left(\theta_i\big(q_i - Q_i^*\big) + \sum_{j\in\mathcal{I}} Q_j^*\right) - \bar{C}_i(q_i) \right\}.$$

Firm i behaves *as though* it believes that if it deviates from Q_i^*, total output will change by the deviation multiplied by the factor θ_i. If $\theta_i = 1$ for all i, we have the Cournot model; if $\theta_i = 0$, then the firms are acting as though their output has no impact on total output, so it is a price-taker. More generally, high conduct parameters are interpreted as firms acting less competitively. A significant literature in empirical IO seeks to measure conduct parameters.

The data set $\mathcal{O}$ is said to be *θ-CV rationalizable* if there are cost functions $\bar{C}_i$ (for all $i \in \mathcal{I}$) and inverse demand functions $\bar{P}_t$ (at each observation t) such that $\bar{P}_t(Q_t) = P_t$ and $(Q_{i,t})_{i\in\mathcal{I}}$ constitutes a θ-CV equilibrium. A straightforward modification of Theorem 2 will give us the following result: $\mathcal{O}$ is θ-CV rationalizable with convex cost functions (where $\theta \gg 0$) if and only if there exists a set of nonnegative real numbers, $\{\delta_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$, that satisfy the co-monotone property and the *generalized common ratio property*

$$(8) \qquad \frac{P_t - \delta_{1,t}}{\theta_1 Q_{1,t}} = \frac{P_t - \delta_{2,t}}{\theta_2 Q_{2,t}} = \cdots = \frac{P_t - \delta_{I,t}}{\theta_I Q_{I,t}} > 0 \quad \text{for all } t \in \mathcal{T}.$$

The most important thing to notice in this result is that *if condition (8) is satisfied by θ, then it is satisfied by λθ for any λ > 0*. This means that θ can only be tested up to scalar multiples, and testing for the absolute level of market power is impossible in our context. However—and this is crucial for our purposes—relative market power is testable. Put another way, our minimal assumptions on costs and demand mean that we could not test the hypothesis that, for a particular firm i, $\theta_i = 1$. However, this does not mean that we could not test the Cournot model, because we could still test the weaker hypothesis that θ_i is the same across firms. The test of the Cournot model we developed in Theorem 2 can be interpreted as a test of the *symmetry* of market interaction as measured by θ_i (for all $i \in \mathcal{I}$). When a data set passes that test, it is consistent with the Cournot hypothesis, but it is also consistent with the θ-CV hypothesis, where $\theta = (\lambda, \lambda, \dots, \lambda)$ for any $\lambda > 0$; in that sense, the conclusion is weak. On the other hand, when a data set fails that test, it is a strong result because all levels of *symmetric* market power have been excluded.

Our observations here are consistent with the results of Bresnahan (1982) and Lau (1982), who showed that the identification of θ requires sufficiently rich variation in (and information on) demand behavior across observations; in contrast, our setup requires no information on the determinants of demand. The following example illustrates the possibility of narrowing down the value of θ when the researcher has access to more information on demand.

EXAMPLE 4: Consider a duopoly with firms i and j where the following observations occur:

(i) At observation t, $P_t = 10$, $Q_{i,t} = 5/3$, $Q_{j,t} = 5/3$, and $d\bar{P}_t/dq \geq -3$.
(ii) At observation t', $P_{t'} = 4$, $Q_{i,t'} = 2$, and $Q_{j,t'} = 5/3$.

We claim that these observations are compatible with $\theta = (3, 3)$, but not with $\theta = (1, 1)$.

Compatibility with $\theta = (3, 3)$ is confirmed if we could find $\delta_{i,t}$, $\delta_{i,t'}$, $\delta_{j,t}$, and $\delta_{j,t'}$ that solves

$$\frac{10 - \delta_{i,t}}{5} = \frac{10 - \delta_{j,t}}{5} \quad \text{and} \quad \frac{4 - \delta_{i,t'}}{6} = \frac{4 - \delta_{j,t'}}{5}$$

and with $\delta_{i,t} \leq \delta_{i,t'}$. It is straightforward to check that these conditions are met if $\delta_{i,t} = 3$, $\delta_{i,t'} = 3$, $\delta_{j,t} = 3$, and $\delta_{j,t'} = 19/6$. In this case, the rationalizing inverse demand function $\bar{P}_t$ can be chosen to satisfy $d\bar{P}_t/dq = -(10 - 3)/5 = -7/5$, which is greater than -3.

Suppose, contrary to our claim, that the data set is Cournot rationalizable with a rationalizing demand $\bar{P}_t$ that satisfies $d\bar{P}_t/dq \geq -3$. The first-order condition of firm i gives

$$\frac{10 - m_{i,t}}{5/3} = -\frac{d\bar{P}_t}{dq} \leq 3,$$

where $m_{i,t}$ is a subgradient of firm i's cost function at output $Q_{i,t} = 5/3$. Therefore, $m_{i,t} \geq 5$, which means that the marginal cost at $Q_{i,t'} = 2$ must be at least 5, since firm i's cost function is convex. However, the price at t' is just 4, so there is a contradiction.

6. APPLICATION: THE WORLD MARKET FOR CRUDE OIL

Accounting for roughly one-third of global oil production, OPEC is a dominant player in the international oil market. OPEC exists, in its own words, "to co-ordinate and unify petroleum policies among Member Countries, in order to secure fair and stable prices for petroleum producers; an efficient, economic and regular supply of petroleum to consuming nations; and a fair return on capital to those investing in the industry." OPEC's stated aims are effectively those of a cartel, but its ability to set world oil prices is questionable and a large literature has emerged that attempts to model its actions. For the most part, the literature suggests that OPEC is a weakly functioning cartel of some sort and is not competitive in either the price-taking or the Cournot sense (see, e.g., Alhajji and Huettner (2000), Dahl and Yücel (1991), Griffin and Neilson (1994), or Smith (2005)). Many of these tests rely on parametric assumptions about the functional forms taken by market demand, countries' objective functions, and production costs. Typically, they also require that factors that shift the cost and inverse demand functions be observed, and rely on constructed proxies such as estimates of countries' extraction costs, the presence of U.S. price controls, and involvement of an OPEC member in a war. Given the ambitious questions they are trying to answer, this seems unavoidable.

Our objective is more specific. All we wish to do is to test whether the behavior of the oil-producing countries is consistent with the Cournot model or, more generally (given the discussion in Section 5), any *symmetric* CV model. Our test gives, so to speak, the greatest benefit of the doubt to the Cournot hypothesis by allowing for a very large class of cost functions and by not making any assumptions at all about the evolution of demand. In spite of this apparent permissiveness, we can reject the Cournot hypothesis in real world data.

Two sources of data are used for this study. The first is the *Monthly Energy Review* (MER), published by the U.S. Energy Information Administration. This provides a full-precision series of monthly crude oil production in thousands of barrels per day by the 12 current OPEC members (Algeria, Angola, Ecuador, Iran, Iraq, Kuwait, Libya, Nigeria, Qatar, Saudi Arabia, the United Arab Emirates, and Venezuela) and 7 nonmembers (Canada, China, Egypt, Mexico, Norway, the United States, and the United Kingdom).[18] This series also contains total world output. The data are available from January 1973 until April 2009, giving a total length of 436 months and $(12 + 7) \times 436 = 8284$ country-month observations. The second source of data is a series of oil prices published by the St. Louis Federal Reserve, in dollars per barrel. This series is deflated by the monthly consumer price index reported by the Bureau of Labor Statistics, so that prices are in 2009 U.S. dollars. Since the time windows over which Cournot behavior is tested are short (12 months or less), the adjustment for inflation should not matter to the results.[19]

We divide the entire data series into multiple subsets, with each set consisting of I countries (where $I = 2, 3, 6$, or 12) and T consecutive months (with $T = 3$, 6, or 12), and then test for Cournot rationalizability on each of these sets. In other words, a data set that contains the output over T consecutive months of I countries corresponds to T observations of the actions in an I-player game and we test whether the data are consistent with a Cournot equilibrium being played at each of the T observations, using the linear program specified in Theorem 2.

The reasons to run the test on subsets of the data as opposed to the entire data set are threefold. First, this demonstrates the ability of the test to reject the Cournot hypothesis on real data that consist of just a small number of observations over a few firms (such as three observations of two firms), in spite of the permissiveness of our nonparametric framework. Second, by considering short windows, we minimize the possibility of spurious rejections due to cost functions changing. Were we to run the test once on the entire data set,

[18]Russia and the former Soviet Union are not used here, because the two are not comparable units. Although the composition of OPEC has changed over the course of the data (Ecuador left in 1994 and returned in 2007, Gabon left in 1995, Angola joined in 2007, and Indonesia left in 2007), the overall pattern of rejecting Cournot behavior below does not depend on what countries are considered to be part of OPEC.

[19]The rejection rates reported below are similar with nominal price series.

TABLE I

REJECTION RATES

OPEC Data Sets

		Number of Countries (I)			
		2	3	6	12
Window (T)	3 Months	0.28	0.54	0.89	1.00
	6 Months	0.65	0.89	1.00	1.00
	12 Months	0.90	0.99	1.00	1.00

Non-OPEC Data Sets

		Number of Countries (I)			
		2	3	6	7
Window (T)	3 Months	0.44	0.75	0.99	1.00
	6 Months	0.83	0.98	1.00	1.00
	12 Months	0.96	1.00	1.00	1.00

Cournot rationalizability will be rejected, but it is possible that this rejection is simply because countries' cost functions have changed significantly over that period. Last, by considering different subsets of countries, we can rule out that the test is rejected solely because a few countries are not responding optimally to the others. In other words, if a particular subset of countries are playing a Cournot game against each other, we will be able to pick up that phenomenon.

Our findings are reported in Table I. (The data and code for the test can found in the Supplemental Material (Carvajal, Deb, Fenske, and Quah (2013)) to this paper.) The rejection rate reported is the proportion of cases that were rejected. For example, there are $436 + 1 - 3 = 434$ 3-month periods in the data and there are 66 possible combinations of 2 out of 12 OPEC members. This means that the entire data series can be divided into $434 \times 66 = 28{,}644$ data sets. We perform our test for Cournot rationalizability on each of these data sets. The entry for two countries and 3 months in Table I reports that out of the 28,644 tests of two OPEC members over 3 months, 8138, or 28%, could not be rationalized. Notice that the rejection rates are increasing in the number of countries and in the number of observations. This must be so, since if a data set consisting of (for example) three countries over 6 consecutive months passes the test, then a subset of this data set involving two of those countries over a shorter window must also pass the test.[20]

[20]Our empirical approach is broadly similar to revealed preference tests of utility maximization based on Afriat's theorem and its variations. In those cases, the test is repeated on multiple data sets and the number of data sets that pass or fail the test, as a proportion of all data sets, is

The results reported in Table I are unambiguous: once there is more than a handful of observations used for the test, the behavior of OPEC members cannot be explained by the Cournot model with convex costs. For nearly 90% of 6-month periods with three countries, the test rejects optimal behavior. Once six countries are included, fewer than one 6-month case in 10 thousand can be rationalized. The same test was performed for the non-OPEC countries (see Table I) and the results are also strongly against the model. For almost all 6-month periods, when at least three countries are considered, the data cannot be rationalized by the Cournot model with convex costs.

APPENDIX

Our proof of Theorem 1 requires the following lemma.

LEMMA 3: *For any generic set of observations* $\mathcal{O} = \{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{T}}$, *there are strictly increasing and* $\mathbf{C}^2$ *cost functions* $\bar{C}_i$ *for all* i *such that* $\{\delta_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ *obeys the common ratio property, where* $\delta_{i,t} = \bar{C}_i'(Q_{i,t})$ *and*

$$P_t q_i - \bar{C}_i(q_i) < P_t Q_{i,t} - \bar{C}_i(Q_{i,t}) \quad \textit{for all } q_i \in [0, Q_{i,t}). \tag{9}$$

PROOF: Suppose that at observation t, $Q_{k,t} \geq Q_{i,t}$ for all i in I. Let $\delta_{k,t}$ be any positive number less than P_t. Define $\beta = (P_t - \delta_{k,t})/Q_{k,t}$ and $\delta_{i,t}$ by

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

In this way, we obtain an array of positive numbers $\{\delta_{i,t}\}_{(i,t)\in\mathcal{I}\times\mathcal{T}}$ that satisfy the common ratio property (1) that is a necessary condition for rationalization. Give each firm i a $\mathbf{C}^1$ marginal cost function with strictly positive values, so that the resulting cost function is strictly increasing and $\mathbf{C}^2$. Furthermore, choose the marginal cost function to guarantee the following two properties: the marginal cost at $Q_{i,t}$ is $\delta_{i,t}$ (this specification is unambiguous because of the genericity assumption) and the resulting cost function $\bar{C}_i$ satisfies (9). We could always choose a marginal cost function that satisfies these two conditions

reported (see, for example, Famulari (1995) and Cherchye, De Rock, and Vermeulen (2011) for tests on observational data, and Harbaugh, Krause, and Berry (2001) and Andreoni and Miller (2002) for tests on experimental data). Typically there will be a *large* number of *small* data sets, with each data set consisting of a small number of demand bundles (observed over time) of a consumer or household. Given that revealed preference tests are discrete pass–fail tests, where it is possible for a test to fail with just a pair of inconsistent observation, this is a sensible approach. To perform a revealed preference test on a large data set, the test ought to be incorporated within an econometric framework that allows for errors and permits the researcher to say something about goodness of fit; see, for example, Varian's (1985) extension of Afriat's theorem in this direction. It is also possible to extend our test of Cournot rationalizability along these lines (see Carvajal et al. (2010)).

because while the first condition fixes the value of marginal cost at observed output levels, we are still free to choose the value of marginal cost *between* the observed output levels $\{Q_{i,t}\}_{t\in\mathcal{T}}$ (though such a marginal cost function will not, in general, be increasing). *Q.E.D.*

PROOF OF THEOREM 1: First, we choose cost functions $\bar{C}_i$ that satisfies the conditions guaranteed by Lemma 3. In particular, condition (9) says that if firm i can sell as much as it likes up to $Q_{i,t}$ at price P_t, then it will choose to sell exactly $Q_{i,t}$. In other words, $Q_{i,t}$ is firm i's best response to $\sum_{j\neq i} Q_{j,t}$ if the inverse demand function at t is $\tilde{P}_t$ defined by $\tilde{P}_t(q) = P_t$ for $q \in [0, Q_t]$ and $\tilde{P}_t(q) = 0$ for $q > Q_t$. So we would have proved the theorem if not for the fact that $\tilde{P}_t$ is not a strictly downward-sloping demand function. It remains for us to construct a strictly downward-sloping demand function at each observation t such that firm i's best response at observation t is $Q_{i,t}$, given that its cost function is $\bar{C}_i$.

For each firm i, we define $g_i(q_i) = k_i(q_i - Q_{i,t}) + \delta_{i,t}$. The graph of g_i is a line, with slope k_i that passes through the point $(Q_{i,t}, \delta_{i,t})$. Since $\delta_{i,t} < P_t$ and $\bar{C}_i$ is $\mathbf{C}^2$, there are $\varepsilon > 0$ and k_i (for $i \in \mathcal{I}$) such that $P_t > g_i(Q_{i,t} - \varepsilon)$, and for q_i in the interval $[Q_{i,t} - \varepsilon, Q_{i,t})$, we have

$$g_i(q_i) > \bar{C}_i'(q_i). \tag{10}$$

(Note that $k_i < 0$ if $\bar{C}_i''(Q_{i,t}) < 0$.) For $q_i \in [0, Q_{i,t} - \varepsilon]$, there exists $\zeta > 0$ such that

$$Pq_i - \bar{C}_i(q_i) < PQ_{i,t} - \bar{C}_i(Q_{i,t}) \quad \text{for} \quad P_t < P < P_t + \zeta; \tag{11}$$

this follows from (9). Note that ζ is common across all firms.

We specify the function $\bar{P}_t'$, so $\bar{P}_t$ can be obtained by integration. Holding the output of firm j (for $j \neq i$) at $Q_{j,t}$, we denote the marginal revenue function for firm i by $\bar{m}_{i,t}$; that is, $\bar{m}_{i,t}(q_i) = \bar{P}_t'(\sum_{j\neq i} Q_{j,t} + q_i)q_i + \bar{P}_t(\sum_{j\neq i} Q_{j,t} + q_i)$. We first consider the construction of $\bar{P}_t'$ in the interval $[0, Q_t]$, where $Q_t = \sum_{i\in\mathcal{I}} Q_{i,t}$. Choose $\bar{P}_t'$ with the following properties: (a) $\bar{P}_t'(Q_t) = (\delta_{i,t} - P_t)/Q_{i,t}$ (which is equivalent to the first-order condition $\bar{m}_{i,t}(Q_{i,t}) = \bar{C}_i'(Q_{i,t}) = \delta_{i,t}$; this is possible because $\{\delta_{i,t}\}_{i\in\mathcal{I}}$ obeys the common ratio property), (b) $\bar{P}_t'$ is negative, decreasing, and concave in $[0, Q_t]$, (c) $\int_0^{Q_t} \bar{P}_t'(q)\,dq = P_t - \bar{P}_t(0) > -\zeta$, and (d) $\bar{P}_t'(Q_t - \varepsilon)$ is sufficiently close to zero so that $\bar{m}_{i,t}(Q_{i,t} - \varepsilon) > g_i(Q_{i,t} - \varepsilon)$. Property (b) guarantees that $\bar{m}_{i,t}$ is decreasing and concave (as a function of q_i). This fact, together with (a) and (d), ensures that $\bar{m}_{i,t}(q_i) > g_i(q_i)$ for all i and q_i in $[Q_{i,t} - \varepsilon, Q_{i,t})$; combining with (10), we obtain $\bar{m}_{i,t}(q_i) > \bar{C}_i'(q)$. Therefore, in the interval $[Q_{i,t} - \varepsilon, Q_{i,t}]$, firm i's profit is maximized at $q_i = Q_{i,t}$. Because of (c), $P_t < \bar{P}_t(q) < P_t + \zeta$, so by (11), $\bar{P}_t(\sum_{j\neq i} Q_{j,t} + q_i)q_i - \bar{C}_i(q_i) < P_t Q_{i,t} - \bar{C}_i(Q_{i,t})$ for q_i in $[0, Q_{i,t} - \varepsilon]$.

To recap, we have constructed $\bar{P}_t'$ (and hence $\bar{P}_t$) such that, with this inverse demand function, firm i's profit at $Q_{i,t}$ is higher than at any output below $Q_{i,t}$, as long as other firms are producing $\sum_{j\neq i} Q_{j,t}$. Our next step is to show how to specify $\bar{P}_t'$ for $q > Q_t$ in such a way that firm i's profit at $q_i = Q_{i,t}$ is higher than at any output level above $Q_{i,t}$ for every firm i. It suffices to have $\bar{P}_t$ such that, for $q_i > Q_{i,t}$,

$$\bar{m}_{i,t}(q_i) = \bar{P}_t'\left(\sum_{j\neq i} Q_{j,t} + q_i\right)q_i + \bar{P}_t\left(\sum_{j\neq i} Q_{j,t} + q_i\right) < \bar{C}_i'(q_i),$$

so firm i's marginal cost always exceeds its marginal revenue for $q_i > Q_{i,t}$. Provided $\bar{P}_t$ is decreasing, it suffices to have $\bar{P}_t'(\sum_{j\neq i} Q_{j,t} + q_i)q_i + P_t < \bar{C}_i'(q_i)$, which is equivalent to

$$-\bar{P}_t'(Q_t + x) > \frac{P_t - \bar{C}_i'(Q_{i,t} + x)}{Q_{i,t} + x} \quad \text{for} \quad x > 0 \text{ and all firms } i. \tag{12}$$

The right side of this inequality is a finite collection of continuous functions of x, and at $x = 0$, the two sides are equal to each other (because of the first-order condition). Clearly, we can choose $\bar{P}_t' < 0$ such that (12) holds for $x > 0$. *Q.E.D.*

PROOF OF THEOREM 3: To see that (A) implies (B), suppose that the data are rationalized with demand functions $\{\bar{P}_t\}_{t\in\mathcal{T}}$ and cost functions $\{\bar{C}_i\}_{i\in\mathcal{I}}$. We have already shown in Theorem 2 that the first-order condition guarantees the existence of $\delta_{i,t} \in \bar{C}_i'(Q_{i,t})$ that obeys the common ratio and co-monotone properties. Since the cost function for firm i is $\mathbf{C}^1$, $\bar{C}_i'(Q_{i,t})$ is unique, so clearly $\delta_{i,t'} = \delta_{i,t}$ whenever $Q_{i,t} = Q_{i,t'}$.

To see that (B) implies (A), it suffices to notice from the proof of Lemma 2 that if the scalars $\{\delta_{i,t}\}_{t\in\mathcal{T}}$ have the fine co-monotone property, $\bar{m}_i$ can be chosen to be a $\mathbf{C}^1$ function. Consequently, $\bar{C}_i$ (as defined by equation (4)) is $\mathbf{C}^2$. *Q.E.D.*

Our proof of Theorem 5 requires the following lemma.

LEMMA 4: *Let $\{[P_t, (Q_{i,t})_{i\in\mathcal{I}}]\}_{t\in\mathcal{S}}$ be a set of $S = |\mathcal{S}|$ observations and let $\{\mu_t\}_{t\in\mathcal{S}}$ be a set of nonnegative real numbers. Let $\Delta_{\mathcal{S}}$ be the set of vectors $\delta = (\delta_{i,t})_{(i,t)\in\mathcal{I}\times\mathcal{S}} \in R_+^{IS}$ that obey the common ratio and co-monotone properties on $\mathcal{S}$ and that satisfy*

$$\frac{P_t - \delta_{i,t}}{Q_{i,t}} \geq \mu_t \quad \textit{for all } i \in \mathcal{I} \textit{ and } t \in \mathcal{S}. \tag{13}$$

Then $\Delta_{\mathcal{S}}$ is a compact and convex sublattice of R_+^{IS}. Furthermore, if $\Delta_{\mathcal{S}}$ is nonempty, it contains its supremum $\delta_{\mathcal{S}}^U$, and $\delta_{\mathcal{S}}^U$ has the following property: for some nonempty set $\mathcal{S}' \subset \mathcal{S}$,

$$\frac{P_t - \delta_{i,t}^U}{Q_{i,t}} = \mu_t \quad \text{for all } i \in \mathcal{I} \text{ and } t \in \mathcal{S}'. \tag{14}$$

PROOF: It is straightforward to show that $\Delta_{\mathcal{S}}$ is a compact and convex sublattice, using broadly the same argument as that used in Theorem 4. Therefore, it is subcomplete and contains its supremum $\delta_{\mathcal{S}}^U$ (see Topkis (1998, Theorem 2.3.1)). Now suppose that for some $\delta \in \Delta_{\mathcal{S}}$, (13) holds with a strict inequality for all $t \in \mathcal{S}$. Then there exists a sufficiently small $\epsilon > 0$ such that $\tilde{\delta}_{i,t} = \delta_{i,t} + \epsilon Q_{i,t}$ obeys the common ratio property, the co-monotone property, and (13). So δ cannot be the supremum of $\Delta_{\mathcal{S}}$. *Q.E.D.*

PROOF OF THEOREM 5: We have already shown in Section 3 that if the UB algorithm terminates affirmatively, then the data set $\mathcal{O}$ is Cournot rationalizable and $\delta^H = (\tilde{\delta}_{i,t}^{\text{ub}})_{(i,t)\in\mathcal{I}\times\mathcal{T}}$. So it remains for us to show that the algorithm will terminate affirmatively at or before T iterations if $\mathcal{O}$ is Cournot rationalizable.

In Step 1, the algorithm sets $\delta_{i,t}^{\text{ub}} = P_t$. If $\mathcal{O}$ is Cournot rationalizable, an application of Lemma 4 (with $\mathcal{S} = \mathcal{T}$, $\mathcal{S}' = \mathcal{T}_1$, and $\mu_t = 0$ for $t \in \mathcal{T}$) tells us that there is a nonempty set $\mathcal{T}_1 \subset \mathcal{T}$ such that for $t \in \mathcal{T}_1$, $\delta_{i,t}^U = P_t$ for all $i \in \mathcal{I}$; therefore, $\delta_{i,t}^{\text{ub}}$ remains unchanged (at P_t, for all $i \in \mathcal{I}$) at every step of the algorithm after Step 1.[21] Notice that the observations in $\mathcal{T}_1$ no longer have a role to play in the algorithm after Step 3, in the sense that the values of $\delta_{i,t}^{\text{ub}}$ and $\gamma_{i,t}^{\text{ub}}$ (for $t \in \mathcal{T} \setminus \mathcal{T}_1$) after that step do not depend on those observations. This is because the upper bound is adjusted downward, and never upward, at each step.

At Step 4, the upper bound is adjusted downward using the common ratio property. Denoting $\mathcal{T} \setminus \mathcal{T}_1$ by $\mathcal{S}$, we claim that there is a nonempty set of observations $\mathcal{T}_2 \subset \mathcal{S}$ such that $\delta_{i,t}^{\text{ub}}$ and $\gamma_{i,t}^{\text{ub}}$ (for all $i \in \mathcal{I}$ and $t \in \mathcal{T}_2$) remain unchanged after that step. Why? Note that the rationalizability of the data set guarantees that the set $\Delta_{\mathcal{S}}$, which consists of vectors $(\delta_{i,t})_{(i,t)\in\mathcal{I}\times\mathcal{S}} \in R_+^{IS}$ satisfying the common ratio and co-monotone properties on $\mathcal{S}$ as well as (13), with $\mu_t = \Lambda_t$, is nonempty. Furthermore, at every step, the vectors $(\delta_{i,t}^{\text{ub}})_{(i,t)\in\mathcal{I}\times\mathcal{S}}$ and $(\gamma_{i,t}^{\text{ub}})_{(i,t)\in\mathcal{I}\times\mathcal{S}}$ generated by the algorithm are upper bounds of $\Delta_{\mathcal{S}}$. Our claim then follows from (14) in Lemma 4.

We claim that there is a nonempty set $\mathcal{T}_3 \subset \mathcal{T} \setminus (\mathcal{T}_1 \cup \mathcal{T}_2)$ such that $\delta_{i,t}^{\text{ub}}$ and $\gamma_{i,t}^{\text{ub}}$ (for all i and $t \in \mathcal{T}_3$) remain unchanged after Step 4 of the second iteration. Clearly, this claim follows from an argument analogous to the one made in the previous paragraph; and a similar claim can be made at Step 4 of the third

[21] We are not saying that $\mathcal{T}_1$ can be *identified* by the algorithm after Step 1, merely that it *exists*.

iteration and so on. Since the sets $\mathcal{T}_1$, $\mathcal{T}_2$, $\mathcal{T}_3, \ldots$ are nonempty, the algorithm cannot continue beyond T iterations. *Q.E.D.*

Afriat's theorem is usually presented as a threefold equivalence of utility maximization, GARP, and the existence of a solution to the Afriat inequalities. Our proof of Theorem 6 below has a similar structure, with MCC playing the role of GARP.

PROOF OF THEOREM 6: We show that the following statements are equivalent:

(A) The set $\mathcal{O}$ is Cournot rationalizable with convex cost functions under common perturbations.

(B) The set $\mathcal{O}$ obeys MCC.

(C) The set $\mathcal{O}$ admits nonnegative numbers $\delta_{i,t}$ and $\phi_{i,t}$ for all (i, t) with the following properties:

(i) For all t,

$$\frac{P_t - \delta_{1,t} - \phi_{1,t}}{Q_{1,t}} = \frac{P_t - \delta_{2,t} - \phi_{2,t}}{Q_{2,t}} = \cdots = \frac{P_t - \delta_{I,t} - \phi_{I,t}}{Q_{I,t}} \geq 0. \tag{15}$$

(ii) We have $\delta_{i,t'} \geq \delta_{i,t'}$ whenever $Q_{i,t'} > Q_{i,t}$.

(iii) We have $\phi_{i,t} \geq \phi_{j,t}$ whenever $Q_{i,t} > Q_{j,t}$.

We have already shown that (A) implies (B). We only sketch the proof that (C) implies (A) since it is straightforward given the proof of Theorem 2. By Lemma 2 and properties (ii) and (iii) in (C), there are convex cost functions $\bar{C}_i$ and f_t with $\delta_{i,t} \in \bar{C}_i'(Q_{i,t})$ and $\phi_{i,t} \in f_t'(Q_{i,t})$. By a straightforward modification of the proof of Lemma 1, we know that at observation t, $Q_{i,t}$ is firm i's best response to $\sum_{j \neq i} Q_{j,t}$ if it has the cost function $\bar{C}_i(q_i) + f_t(q_i)$ and if the inverse demand curve is linear, passes through (Q_t, P_t), and has a slope equal to $-(P_t - \delta_{i,t} - \phi_{i,t})/Q_{i,t}$ (which, by (15), is independent of i).

It remains for us to show that (B) implies (C). The MCC guarantees that there is a complete, reflexive, and transitive order $\succeq^*$ on $\mathcal{I} \times \mathcal{T}$ that agrees with $\succeq$ and $\succ$ in the following sense: if $(i, t) \succeq (j, t')$, then $(i, t) \succeq^* (j, t')$, and if $(i, t) \succ (j, t')$, then $(i, t) \succeq^* (j, t')$ but $(j, t) \not\succeq^* (j, t')$.[22] We write $(i, t) \succ^* (j, t')$ if $(i, t) \succeq^* (j, t')$ but $(j, t) \not\succeq^* (j, t')$. Let $\delta_{i,t}$ be a representation of $\succeq^*$, that is, $\delta_{i,t} \geq (>)\ \delta_{j,t'}$ if $(i, t) \succeq^* (\succ^*)\ (j, t')$, with the property that $0 < \delta_{i,t} < P_t$ for all $(i, t) \in \mathcal{I} \times \mathcal{T}$. Clearly, such a representation exists and will obey property (ii). For each $t \in \mathcal{T}$, choose $b_t > 0$ and sufficiently small so that $P_t - b_t Q_{i,t} - \delta_{i,t} > 0$ for all $i \in \mathcal{I}$, and

$$P_t - b_t Q_{i,t} - \delta_{i,t} > (=)\ P_t - b_t Q_{j,t} - \delta_{j,t} \tag{16}$$

[22]This claim is almost obvious here since $\mathcal{I} \times \mathcal{T}$ is finite. Essentially the same result in the case where the domain is infinite can be proved using Szpilrajn's theorem (see Ok (2007, Chapter 1, Corollary 1)).

whenever $Q_{i,t} > (=)\ Q_{j,t}$. This is possible because $\delta_{i,t} < (=)\ \delta_{j,t}$ if $Q_{i,t} > (=)\ Q_{j,t}$. Define $\phi_{i,t} = P_t - b_t Q_{i,t} - \delta_{i,t}$; note that $\phi_{i,t} > 0$ and it obeys (iii) because of (16). Furthermore,

$$\frac{P_t - \delta_{i,t} - \phi_{i,t}}{Q_{i,t}} = b_t \quad \text{for all } i \in \mathcal{I},$$

so we obtain (15) (which is property (i)).[23] *Q.E.D.*

REFERENCES

AFRIAT, S. (1967): "The Construction of a Utility Function From Expenditure Data," *International Economic Review*, 8, 67–77. [2351]

ALHAJJI, A., AND D. HUETTNER (2000): "OPEC and World Crude Oil Markets From 1973 to 1994: Cartel, Oligopoly, or Competitive?" *The Energy Journal*, 21 (3), 31–60. [2370]

ANDREONI, J., AND J. MILLER (2002): "Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism," *Econometrica*, 70, 737–753. [2373]

BRESNAHAN, T. (1982): "The Oligopoly Solution Concept Is Identified," *Economics Letters*, 10, 87–92. [2369]

——— (1989): "Empirical Studies of Industries With Market Power," in *Handbook of Industrial Economics*, Vol. II, ed. by R. Schmalensee and D. Willig. Amsterdam: Elsevier. [2368]

BROWN, D., AND R. MATZKIN (1996): "Testable Restrictions on the Equilibrium Manifold," *Econometrica*, 64 (6), 1249–1262. [2351,2352]

CARVAJAL, A., R. DEB, J. FENSKE, AND J. K.-H. QUAH (2010): "Revealed Preferences Tests of the Cournot Model," Discussion Paper 506, Economics Department, Oxford University. [2366, 2373]

——— (2013): "Supplement to 'Revealed Preference Tests of the Cournot Model'," *Econometrica Supplemental Material*, 81, http://www.econometricsociety.org/ecta/supmat/10175_data_and_programs.zip. [2372]

CHERCHYE, L., B. DE ROCK, AND F. VERMEULEN (2011): "The Revealed Preference Approach to Collective Consumption Behaviour: Testing and Sharing Rule Recovery," *Review of Economic Studies*, 78, 176–198. [2373]

DAHL, C., AND M. YÜCEL (1991): "Testing Alternative Hypotheses of Oil Producer Behavior," *The Energy Journal*, 12 (4), 117–138. [2370]

FAMULARI, M. (1995): "A Household-Based, Nonparametric Test of Demand Theory," *Review of Economics and Statistics*, 77 (2), 372–382. [2373]

FORGES, F., AND E. MINELLI (2009): "Afriat's Theorem for General Budget Sets," *Journal of Economic Theory*, 144, 135–145. [2353]

GALAMBOS, A. (2011): "The Complexity of Nash Rationalizability," Report. [2354]

GRIFFIN, J., AND W. NEILSON (1994): "The 1985–86 Oil Price Collapse and Afterwards: What Does Game Theory Add?" *Economic Inquiry*, 32 (4), 543–561. [2370]

HARBAUGH, W. T., K. KRAUSE, AND T. R. BERRY (2001): "GARP for Kids: On the Development of Rational Choice Behavior," *American Economic Review*, 91 (5), 1539–1545. [2373]

[23] We have shown that (B) guarantees that $\delta_{i,t}$ can be chosen to be positive (not just nonnegative) and, similarly, $\phi_{i,t}$ can be chosen to be positive. So in fact the rationalizing cost function in (A) (see (4)) can be chosen to be *strictly* increasing because both $\bar{C}_i$ and f_t can be chosen to be strictly increasing.

KUBLER, F. (2003): "Observable Restrictions of General Equilibrium Models With Financial Markets," *Journal of Economic Theory*, 110, 137–153. [2352]

LAU, L. (1982): "On Identifying the Degree of Competitiveness From Industry Price and Output Data," *Economics Letters*, 10, 93–99. [2369]

OK, E. (2007): *Real Analysis With Economic Applications*. Princeton, NJ: Princeton University Press. [2377]

RAY, I., AND L. ZHOU (2001): "Game Theory via Revealed Preferences," *Games and Economic Behavior*, 37 (2), 415–424. [2354]

ROSSE, J. (1970): "Estimating Cost Function Parameters Without Using Cost Data: Illustrated Methodology," *Econometrica*, 38 (2), 256–275. [2361]

SMITH, J. L. (2005): "Inscrutable OPEC: Behavioral Tests of the Cartel Hypothesis," *The Energy Journal*, 26 (1), 51–82. [2370]

SPRUMONT, Y. (2000): "On the Testable Implications of Collective Choice Theories," *Journal of Economic Theory*, 93, 205–232. [2354]

SZIDAROVSKY, F., AND S. YAKOWITZ (1977): "A New Proof of the Existence and Uniqueness of the Cournot Equilibrium," *International Economic Review*, 18 (3), 787–789. [2360]

TOPKIS, D. (1998): *Supermodularity and Complementarity*. Princeton, NJ: Princeton University Press. [2362,2376]

VARIAN, H. (1985): "Nonparametric Analysis of Optimizing Behavior With Measurement Error," *Journal of Econometrics*, 30, 445–458. [2373]

——— (1988): "Revealed Preference With a Subset of Goods," *Journal of Economic Theory*, 46, 179–185. [2359]

Dept. of Economics, University of Western Ontario, Social Science Centre Room 4071, London, Ontario, N6A 5C2, Canada; acarvaj@uwo.ca,

Dept. of Economics, University of Toronto, 150, St George Street, Toronto, Ontario, M5S 3G7, Canada; rahul.deb@utoronto.ca,

Dept. of Economics, Oxford University, Manor Road Building, Manor Road, Oxford, OX1 3UQ, United Kingdom; james.fenske@economics.ox.ac.uk,

and

Dept. of Economics, Oxford University, Manor Road Building, Manor Road, Oxford, OX1 3UQ, United Kingdom; john.quah@economics.ox.ac.uk.

Manuscript received July, 2011; final revision received April, 2013.