

J.E. Harrington · B.F. Hobbs · J.S. Pang · A. Liu · G. Roch

Collusive game solutions via optimization

It is with great honor that we dedicate this paper to Professor Terry Rockafellar on the occasion of his 70th birthday. Our work provides another example showing how Terry's fundamental contributions to convex and variational analysis have impacted the computational solution of applied game problems.

Received: February 14, 2004 / Accepted: February 16, 2005

Published online: July 14, 2005 – © Springer-Verlag 2005

Abstract. A Nash-based collusive game among a finite set of players is one in which the players coordinate in order for each to gain higher payoffs than those prescribed by the Nash equilibrium solution. In this paper, we study the optimization problem of such a collusive game in which the players collectively maximize the Nash bargaining objective subject to a set of incentive compatibility constraints. We present a smooth reformulation of this optimization problem in terms of a nonlinear complementarity problem. We establish the convexity of the optimization problem in the case where each player's strategy set is unidimensional. In the multivariate case, we propose upper and lower bounding procedures for the collusive optimization problem and establish convergence properties of these procedures. Computational results with these procedures for solving some test problems are reported.

1. Introduction

In industries characterized by repeated interaction, tacit collusion among producers can emerge that enables price to drift above single-period Nash equilibrium levels. An example of such an industry is the restructured electric power generation sector, where auctions are held hourly or half-hourly in many markets [21]. Empirical evidence from the California and England-Wales markets indicate that prices have exceeded single-period Nash prices for some periods of time [19, 23]. Models of tacit collusion can be useful to understand how changed market design or structure might affect prices in these circumstances.

J.E. Harrington: Department of Economics, The Johns Hopkins University, Baltimore, Maryland 21218-2682, USA. e-mail: joe.harrington@jhu.edu.

B.F. Hobbs: Department of Geography and Environmental Engineering, The Johns Hopkins University, Baltimore, Maryland 21218-2682, USA. e-mail: bhobbs@jhu.edu.

This author's research was partially supported by the National Science Foundation under grant ECS-0080577.

J.S. Pang: Department of Mathematical Sciences, Rensselaer Polytechnic Institute, Troy, New York, 12180-3590, USA. e-mail: pangj@rpi.edu.

This author's research was partially supported by the National Science Foundation under grant CCR-0098013.

A. Liu: Department of Applied Mathematics and Statistics, The Johns Hopkins University, Baltimore, Maryland 21218-2682, USA. e-mail: lliu@jhu.edu.

G. Roch: Department of Applied Mathematics and Statistics, The Johns Hopkins University, Baltimore, Maryland 21218-2682, USA. e-mail: grantroch@jhu.edu.

Mathematics Subject Classification (1991): 90C26, 90C33, 90C90, 91A10, 91A20

Broadly speaking, there are two general approaches to modeling tacit equilibria. One is agent-based simulation, in which autonomous agents learn and evolve strategies that can mimic tacit collusion. This approach has been applied, for example, to the analysis of market power in the England-Wales energy generation market [3]. The other approach is dynamic equilibrium models, or “supergames”, of the type which is the main focus of this paper. In general, collusion involves firms coordinating their quantity and price decisions for the purpose of generating higher payoffs. In considering the incentives to collude, it is well known that equilibria for most oligopoly models are Pareto-inefficient; that is, all firms could increase their payoff by jointly modifying their decisions. (For a general statement about the Pareto-inefficiency of Nash equilibria, see Dubey [6].) This Pareto-inefficiency takes the form that all firms would realize a higher payoff if they marginally reduced their quantities. The approach in the economics literature to modelling collusion is to enrich the game by having firms make decisions repeatedly. Quantities that are not equilibria when firms interact only once can be equilibria in the dynamic setting. Of course, if the quantities are not Nash equilibria for the static game, firms can increase their instantaneous payoff by producing differently. To offset this short-run gain, dynamic strategies have firms respond to any such deviation by acting in a manner to reduce a deviating firm’s future payoff. This defines a new set of equilibrium conditions that expands the set of equilibrium quantities as one moves from the static to the dynamic game. For the purpose of the ensuing discussion, let Ω denote the set of equilibrium quantities for the dynamic game.

There is one serious weakness with the movement from the static to the dynamic game – the loss of uniqueness of equilibrium. Under restrictive but plausible assumptions, the static Cournot game has a unique Nash equilibrium. In contrast, the dynamic Cournot game generally has many quantities consistent with Nash equilibrium. This leaves open the issue of selecting an element from Ω . When firms are symmetric, it has been common practice in the economics literature to focus on the best symmetric element of Ω . However, when firms are asymmetric, such as in their cost and capacities, there is no such focal point. The selection from Ω should be asymmetric but how exactly should it relate to firms’ traits?

In thinking about this problem from the firms’ perspective, they are likely to disagree over which element of Ω to choose; each firm wanting to select the one that gives it the highest payoff. Firms with lower costs probably prefer lower prices and, given any price, each firm desires a bigger market share. In light of such disagreement, it is natural to think of firms bargaining to achieve some resolution. This is the basis for the selection approach formulated in Harrington [12] which uses the set of equilibrium quantities to construct a bargaining problem. In the axiomatic bargaining literature (see, e.g., Osborne and Rubinstein [18]), a bargaining problem is defined by a set of payoff vectors, Γ , which is the set over which players bargain, and a disagreement payoff, $d \in \Gamma$, which is the payoff vector if they fail to reach an agreement. One then applies a bargaining solution to this problem which, under certain conditions, produces a unique solution. Applied to our setting, the approach of Harrington [12] is to specify Γ to be the payoff vectors induced by quantities in Ω and d to be the Nash equilibrium payoff for the static game. If one uses the bargaining solution of Nash [17], the collusive problem can be represented by the following optimization problem:

$$\max_{q \in \Omega} \prod_{f \in \mathcal{F}} \left(\pi_f(q) - \pi_f^N \right), \quad (1)$$

where q is the vector of firms' quantities, $\pi_f(q)$ is the payoff of firm f , π_f^N is the static Nash equilibrium payoff for firm f , and \mathcal{F} is the set of firms. To provide some perspective of the formulation (1), previous work had formulated the selection problem associated with collusion as a bargaining problem but had specified the choice set to be the set of feasible quantities, which we will denote X , rather than the set of equilibrium quantities; see, for example, Schmalensee [22]. It is now well known that such an approach is flawed because if the solution fails to lie in Ω then firms have agreed to an outcome that they have no intent of implementing. Though it is then well motivated to replace X with Ω , we are replacing the (typically) convex set X with a set that may not be convex. Based on their descriptive appeal, we focus on pure stationary outcomes, which prevent achieving convexification through randomization or other means.

Besides providing a formal motivation, this paper aims at studying the collusive optimization problem (1) from a global optimization perspective. Due to the possible nonconcavity of the objective function and the nonconvexity of constraint set, the stated aim is computationally challenging. Our main contributions consist of a set of supporting results that provide mathematical insights into the problem, the development of upper and lower bounding procedures for the global solution of (1), and the demonstration of convergence properties of these procedures.

The organization of the rest of the paper is as follows. In the next section, we present a summary of the variational inequality (VI) approach to computing noncooperative Nash equilibria. This is followed by two sections in which we define and analyze the feasible set and objective function, respectively, of (1). In particular, we establish a convexity result in the univariate case in which the strategy set of each player of the game is a compact interval; convexity of the collusive game optimization problem is established in this case. In Section 5, we present upper and lower bounding procedures in the multivariate case and establish some convergence results for these procedures. Finally, in Section 6, we illustrate the univariate and multivariate cases with some numerical results.

2. The VI approach to computing Nash equilibria

We begin with a review of the VI approach to the well-known noncooperative Nash equilibrium problem. For a comprehensive study of finite-dimensional variational inequalities, see [7]. The players of this game are labelled by the elements f in a finite index set \mathcal{F} (for firms). Player f 's strategy set is denoted X_f , which is a nonempty, convex, and compact subset of the Euclidean space \mathfrak{R}^{n_f} for some positive integer n_f . Elements of X_f are denoted by q_f , which are n_f -dimensional vectors. Let

$$X \equiv \prod_{f \in \mathcal{F}} X_f \subseteq \mathfrak{R}^n, \quad \text{where } n \equiv \sum_{f \in \mathcal{F}} n_f.$$

Elements of X are denoted by q , whose components are q_f for all $f \in \mathcal{F}$. We write

$$X_{-f} \equiv \prod_{t \in \mathcal{F}, t \neq f} X_t, \quad \forall f \in \mathcal{F},$$

and write q_{-f} to denote an arbitrary element of X_{-f} ; thus q_{-f} is the vector with components q_t for all $t \neq f$. Player f 's payoff function is denoted π_f , which is a real-valued function defined on \mathfrak{R}^n ; thus the payoff $\pi_f(q)$ to player f depends on the vector of all players' strategies. We use the alternative notation $\pi_f(q_f, q_{-f})$ for $\pi_f(q)$ when we want to highlight the dependence of player f 's payoff on his own strategy q_f and his rivals' collective strategy q_{-f} . Throughout the paper, we postulate that $\pi_f(\cdot, q_{-f})$ is a strictly concave function in the argument q_f for each fixed but arbitrary $q_{-f} \in X_{-f}$ and $\pi_f(q_f, \cdot)$ is convex in the argument q_{-f} for each fixed but arbitrary $q_f \in X_f$.

Parameterized by q_{-f} , player f 's payoff maximization problem is the optimization problem in the primary variable q_f :

$$\begin{aligned} &\text{maximize } \pi_f(q_f, q_{-f}) \\ &\text{subject to } q_f \in X_f. \end{aligned} \tag{2}$$

A Nash equilibrium is a tuple $q^N \equiv (q_f^N : f \in \mathcal{F})$ such that for all $f \in \mathcal{F}$, $q_f^N \in X_f$ and

$$\pi_f(q^N) \geq \pi_f(q_f, q_{-f}^N), \quad \forall q_f \in X_f.$$

We write $\pi_f^N \equiv \pi_f(q^N)$ to denote the firms' Nash payoffs. Define the vector function

$$F(q) \equiv -(\nabla_{q_f} \pi_f(q) : f \in \mathcal{F}), \quad q \in X. \tag{3}$$

For ease of later reference, we summarize in the following result properties of the optimization problem (2) and the Nash equilibrium.

Proposition 1. *Let X_f be a nonempty compact convex subset of \mathfrak{R}^{n_f} ; let $\pi_f : \mathfrak{R}^n \rightarrow \mathfrak{R}$ be continuously differentiable and such that $\pi_f(\cdot, q_{-f})$ is strictly concave for every $q_{-f} \in X_{-f}$. The following statements hold.*

(a) *For all $q_{-f} \in X_{-f}$, a unique maximizer, denoted $q_f^*(q_{-f})$, exists, which satisfies*

$$(q_f - q_f^*(q_{-f}))^T \nabla_{q_f} \pi_f(q_f^*(q_{-f}), q_{-f}) \leq 0, \quad \forall q_f \in X_f.$$

(b) *If $\pi_f(q_f, \cdot)$ is convex for every $q_f \in X_f$ then the optimal value function*

$$\pi_f^*(q_{-f}) \equiv \pi_f(q_f^*(q_{-f}), q_{-f})$$

is convex and continuously differentiable with

$$\nabla \pi_f^*(q_{-f}) \equiv \nabla_{q_{-f}} \pi_f(q_f^*(q_{-f}), q_{-f}).$$

(c) *If X_f is polyhedral, then $q_f^*(q_{-f})$, and thus $\nabla \pi_f^*(q_{-f})$, are piecewise smooth functions of their argument.*

(d) *A Nash equilibrium q^N exists; moreover q^N satisfies*

$$F(q^N)^T (q - q^N) \geq 0, \quad \forall q \in X.$$

(e) *If $F(q)$ is a strictly monotone function on X , then the Nash equilibrium q^N is unique.*

Proof. The existence, uniqueness, and variational characterization of $q_f^*(q_{-f})$ are standard results in convex programming. The convexity and continuous differentiability of the value function $\pi_f^*(q_{-f})$ are immediate consequences of the well-known Danskin's Theorem [4]. The assertion about the piecewise smoothness of $q_f^*(q_{-f})$ and $\nabla\pi_f^*(q_{-f})$ follows from sensitivity results of parametric nonlinear programs/variational inequalities; see [2, 7, 20]. Finally, the assertions about the Nash equilibrium are well-known results; see e.g. [7]. \square

In the rest of this paper, we assume that the Nash equilibrium q^N is unique; see Proposition 2 where this assumption is used.

2.1. The firms' payoff functions

For the most part in this paper, we are interested in a Cournot oligopoly problem in which the payoff function $\pi_f(q)$ is of the following form:

$$\pi_f(q) \equiv p(Q)^T q_f - c_f(q_f), \quad (4)$$

where $p(Q)$ is the market inverse demand (vector) function, with

$$Q \equiv \sum_{f \in \mathcal{F}} q_f$$

being the industry production, i.e., total production by all firms (i.e., players), and $c_f(q_f)$ is firm f 's production cost function. We write $Q_{-f} \equiv Q - q_f$ for the total production by firm f 's rivals. Associated with the Nash equilibrium q^N , we write Q^N and Q_{-f}^N for the industry production and the industry less firm f 's production, respectively. The form of the payoff function (4) implies that n_f is a constant N for all firms. This is *not* a restrictive assumption. Indeed, suppose that each firm has n_f plants, with n_f not necessarily equal. Then define $N = \max_{f \in \mathcal{F}} n_f$. We can add $N - n_f$ additional plants to firm f 's collection of plants, each with zero capacity. This minor adjustment will not alter the most difficult technical challenge of the problem (1), which is its global optimality and the main concern of our work.

Both functions $p(Q)$ and $c_f(q_f)$ in (4) are assumed to be continuously differentiable, and additionally, c_f is assumed to be convex and nonnegative. The algorithmic treatment in this paper pertains to the further special case where $p(Q)$ is a separable affine function:

$$p(Q) \equiv \alpha - \text{Diag}(\beta) Q, \quad (5)$$

where α and β are positive vectors and $\text{Diag}(\beta)$ is the diagonal matrix whose diagonal entries are the components of β . For the most part, the cost function $c_f(q_f)$ is not required to be linear, except in Subsection 5.2 and in the computational tests. With a separable affine $p(Q)$ as above, it is easy to see that the vector function $F(q)$ in (3) is strictly monotone with a positive definite Jacobian matrix; consequently, the Nash equilibrium q^N is naturally unique in a Cournot oligopoly problem.

When dealing with the payoff function (4), we employ the following notation: $q_f^*(Q_{-f})$ for $q_f^*(q_{-f})$, $\pi_f^*(Q_{-f})$ for $\pi_f^*(q_{-f})$, and $\pi_f(q_f, Q_{-f})$ for $\pi_f(q)$; this change of notation is more descriptive because player f 's optimal response and his payoff in this case depend on his rivals' strategies only through the cumulative output Q_{-f} and not directly through the individual quantities q_{-f} .

3. The family of collusive sets

This section introduces the family of collusive sets Ω_δ for $\delta \in [0, 1]$. We begin by motivating these sets using an infinitely repeated extension of the basic Nash game. Some properties of elements of Ω_δ are then derived. We conclude the section by defining the "Pareto improvement" strategies, which are central to the Nash bargaining optimization problem to be formulated and discussed in the next section.

3.1. Motivation

Although noncooperative Nash equilibria involve each firm individually maximizing its payoff, it is generally not true that firms are collectively maximizing some objective such as the sum of their payoffs. As mentioned in the Introduction, Nash equilibria are generally not Pareto-efficient. In many formulations of the Cournot game, this Pareto-inefficiency takes the form that all firms would have a higher payoff if they all marginally reduced their quantities. It is the possibility that all firms could benefit from coordination of their quantity choices that provides the basis for observed episodes of collusion in actual markets. As such coordination does not emerge as equilibrium behavior in a standard Cournot game, generation of such behavior requires either forsaking the assumption of equilibrium or altering the specification of the game. In the economics literature [10], it is typical to maintain equilibrium as a description of firm behavior and instead modify the game in the direction of greater descriptive realism. In particular, let us now assume that firms choose quantities repeatedly rather than only once.

The infinitely repeated extension of our game has firms choose quantities and realize payoffs in each of an infinite number of periods. (What is critical for the ensuing analysis is not that there is an infinite number of periods but rather that there is no upper bound on the number of periods. All results could be derived assuming that the number of periods is stochastic as long as, in all periods, a firm assigns positive probability to there being at least one more period. For a discussion of repeated games, see Fudenberg and Tirole [10].) In defining the strategy space of the infinitely repeated extension, assume that a firm in period t knows the quantities selected over the previous $t - 1$ periods so that its strategy can condition on that information. A strategy for firm f in the infinitely repeated extension is then an infinite sequence of functions, $\rho_f \equiv \left\{ \rho_f^t \right\}_{t=1}^\infty$, where $\rho_f^t : X^{t-1} \rightarrow X_f$ and X^{t-1} represents the space of period t histories. The strategy set of a firm is the space of all infinite sequences of such functions. Let ρ denote the vector of firms' strategies and $q^{t-1} \in X^{t-1}$ represent a period t history. A firm's payoff function for the infinitely repeated extension is assumed to be the sum of discounted single-period payoffs,

$$\sum_{t=1}^{\infty} \delta^{t-1} \pi_f \left(\rho^t \left(q^{t-1} \right) \right),$$

where q^{t-1} is defined recursively by ρ ; that is, it is the sequence of quantity vectors induced by firms' strategies. The scalar $\delta \in [0, 1]$ is a firm's discount factor and measures how a firm values a single-period payoff in the next period compared to one in the current period. The weight given to a future single-period payoff declines geometrically in the number of periods. Since $\delta \leq 1$ then current single-period payoffs are valued at least as much as those in the future. Also note that as δ declines, less weight is given to future single-period payoffs and when $\delta = 0$ firms care only about their current single-period payoff.

Even if there is a unique Nash equilibrium for the single-period game, it is well known that the set of Nash equilibria for its infinitely repeated extension can be large. It is assured of being nonempty since it includes the infinite repetition of the Nash equilibrium for the single-period game: firm f produces q_f^N in every period for all f (so that there is no conditioning on the history). More interesting is to consider Nash equilibria that result in quantity vectors distinct from q^N and, in particular, generate higher payoffs. This returns us to the issue of collusion which in the context of the infinitely repeated extension means an equilibrium that yields a higher average payoff than π_f^N , for all f . (The central result in the theory of repeated games is the Folk Theorem. Define v_f as firm f 's minimax payoff in the single-period game. If $\bar{\pi}_f > v_f \forall f$ and $\exists q \in X$ such that $\pi_f(q) = \bar{\pi}_f$ then there exists a Nash equilibrium for the infinitely repeated extension such that firm f 's average single-period payoff is $\bar{\pi}_f \forall f$, when δ is sufficiently close to one. Details can be found in Fudenberg and Tirole (1991).) Arguably the simplest strategy achieving that objective is the grim trigger strategy [9]. It is defined as follows: for all $f \in \mathcal{F}$,

$$\begin{aligned} \rho_f^1 &= q_f \\ \rho_f^t &= \left\{ \begin{array}{ll} q_f & \text{if } q^\tau = q \quad \forall \tau \in \{1, 2, \dots, t-1, \} \\ q_f^N & \text{otherwise;} \end{array} \right\} \quad \forall t \in \{2, 3, \dots\}, \quad (6) \end{aligned}$$

where $q \in X$ and is to be interpreted as the collusive quantity vector. This strategy says that firm f chooses q_f in period 1. In any future period, it chooses q_f if all past quantity vectors have been q . Given each of the other firms deploys (6), the payoff to firm f from doing so is $\pi_f(q)/(1 - \delta)$ as this strategy profile results in it receiving profit of $\pi_f(q)$ in every period. For (6) to be a Nash equilibrium, $\pi_f(q)/(1 - \delta)$ must be at least as great as that earned from using any other strategy. Another strategy yields a different payoff for firm f only if it entails firm f producing a quantity different from q_f in some period. Without loss of generality, consider a strategy that does so in the first period. Note that the future payoff is $\pi_f^N/(1 - \delta)$ regardless of what firm f produces as long as it is different from q_f . This follows from each of the other firms using (6) and that the best response of firm f is to produce q_f^N . A necessary and sufficient condition for the

strategy (6) to be a Nash equilibrium is then: for all $f \in \mathcal{F}$,

$$\begin{aligned} \frac{\pi_f(q)}{1-\delta} &\geq \pi_f(q'_f, q_{-f}) + \delta \frac{\pi_f^N}{1-\delta} \quad \forall q'_f \neq q_f \\ \Leftrightarrow \frac{\pi_f(q)}{1-\delta} &\geq \pi_f^*(q_{-f}) + \delta \frac{\pi_f^N}{1-\delta} \\ \Leftrightarrow \pi_f(q) &\geq (1-\delta)\pi_f^*(q_{-f}) + \delta\pi_f^N. \end{aligned}$$

In the repeated game literature, this condition is called an incentive compatibility constraint. It simply states that unilateral deviation from a collusive solution (“cheating”) should not be more profitable for any firm than continued collusion.

We then define the set of stationary quantity vectors supportable by Nash equilibria to be

$$\Omega_\delta = \{q \in X : \pi_f(q) \geq (1-\delta)\pi_f^*(q_{-f}) + \delta\pi_f^N, \forall f \in \mathcal{F}\}.$$

The logic behind Ω_δ possibly containing quantity vectors different from q^N is as follows. If $q \neq q^N$ then, by instead choosing $q_f^*(q_{-f})$, a firm can raise its current single-period payoff. Counterbalancing this temptation is the reaction of the other firms’ future quantities to such a departure. As long as the inequality in Ω_δ holds, i.e.,

$$\pi_f(q) \geq (1-\delta)\pi_f^*(q_{-f}) + \delta\pi_f^N, \quad (7)$$

then other firms’ future response sufficiently depresses firm f ’s future payoffs that it dominates the gain in its current single-period payoff. To better illustrate this situation, consider the case where the firms’ quantity vectors are scalars. In this case, under the assumption that firm f ’s single-period payoff function is decreasing in q_{-f} (which is standard in the literature and will be substantiated in a Cournot oligopoly model; cf. the proof of Proposition 4), satisfaction of (7) then requires that $q_{-f} < q_{-f}^N$ so that firm f is deterred from producing $q_f^*(q_{-f})$ by the threat that other firms will raise their quantities from q_{-f} to q_{-f}^N in future periods.

In general, there are a number of technical questions associated with the set Ω_δ , such as its nonemptiness, relations between its elements q and the Nash strategy q^N , relations between the associated payoffs $\pi_f(q)$ and the Nash payoffs π_f^N , and the dependence on δ . The following simple result addresses these questions.

Proposition 2. *Let $\delta \in [0, 1]$ be arbitrary. Assume that the Nash equilibrium q^N is unique. For all $q \in \Omega_\delta$, it holds that*

- (a) $\pi_f^*(q_{-f}) - \pi_f^N \geq \pi_f(q) - \pi_f^N \geq (1-\delta)(\pi_f^*(q_{-f}) - \pi_f^N) \geq 0$;
- (b) $\pi_f(q) > (1-\delta)\pi_f^*(q_{-f}) + \delta\pi_f^N \Rightarrow \pi_f(q) > \pi_f^N$;
- (c) *except for $\delta = 1$,*

$$\pi_f^*(q_{-f}) > \pi_f^N \Rightarrow \pi_f(q) > \pi_f^N.$$

Moreover, $\Omega_0 = \{q^N\}$ and

$$\Omega_1 = \left\{ q \equiv (q_f) \in X : \pi_f(q_f, q_{-f}) \geq \pi_f^N, \forall f \in \mathcal{F} \right\}.$$

Finally, for all $0 \leq \delta_1 \leq \delta_2 \leq 1$, $\Omega_{\delta_1} \subseteq \Omega_{\delta_2}$.

Proof. We first note that the two explicit expressions for Ω_0 and Ω_1 are obvious; the former is due to the uniqueness of the Nash equilibrium. The first two inequalities in (a) require no proof. For the third inequality, note that

$$\delta (\pi_f(q) - \pi_f^N) \geq (1 - \delta) (\pi_f^*(q_{-f}) - \pi_f(q)) \geq 0.$$

Hence $(\pi_f(q) - \pi_f^N) \geq 0$ if $\delta > 0$. The same must also hold for $\delta = 0$. Hence (a) holds, from which (b) follows readily, and so does (c).

To prove the last assertion of the proposition, let δ_1 and δ_2 be as given. Let $q \in \Omega_{\delta_1}$. We have

$$\begin{aligned} \pi_f(q_f, Q_{-f}) &\geq \pi_f^*(Q_{-f}) + \delta_1 (\pi_f^N - \pi_f^*(Q_{-f})) \\ &\geq \pi_f^*(Q_{-f}) + \delta_2 (\pi_f^N - \pi_f^*(Q_{-f})), \end{aligned}$$

where the second inequality is by the fact that $\pi_f^N \leq \pi_f^*(Q_{-f})$. Thus $q \in \Omega_{\delta_2}$. \square

The above result shows that the family of collusive sets $\{\Omega_\delta : \delta \in [0, 1]\}$ is expanding from the singleton Ω_0 , which consists of the single element of the Nash equilibrium solution q^N , to the set Ω_1 , which consists of all admissible productions $q \in X$ for which each firm's payoff $\pi_f(q)$ is not lower than the Nash payoff π_f^N .

3.2. Pareto improvements

A basic motivation in considering collusive strategies is to allow firms to earn higher payoffs than their respective Nash payoffs. Mathematically, this raises the question of whether there exist $q \in \Omega_\delta$ such that $\pi_f(q) > \pi_f^N$ for all $f \in \mathcal{F}$. We call such a quantity vector q a *Pareto improvement*. Due to the fundamental role of this question in the notion of collusion, we give a sufficient condition for a positive answer to the question, based on the familiar concept of feasible ascent in optimization.

Proposition 3. *If there exists $q \in X$ such that*

$$\nabla \psi_{\delta, f}(q^N)^T (q - q^N) > 0, \quad \forall f \in \mathcal{F}, \quad (8)$$

where $\psi_{\delta, f}(q) \equiv \pi_f(q) - (1 - \delta)\pi_f^*(q_{-f})$, then a Pareto improvement exists.

Proof. It is easy to see that, for q satisfying (8), the vector $q(\tau) \equiv q^N + \tau(q - q^N)$ is a Pareto improvement for all $\tau > 0$ sufficiently small. Indeed, for such a τ , we have

$$\psi_{\delta, f}(q(\tau)) > \psi_{\delta, f}(q^N) = \delta \pi_f^N$$

for all $f \in \mathcal{F}$. By part (b) of Proposition 2, it follows that $\pi_f(q(\tau)) > \pi_f^N$. \square

It is interesting to specialize (8) to the payoff function (4). Noting that

$$\nabla_{q_t} \pi_f(q) = \begin{cases} p(Q) + Jp(Q)^T q_f - \nabla c_f(q_f) & \text{if } t = f \\ Jp(Q)^T q_f & \text{if } t \neq f, \end{cases}$$

where $Jp(Q)$ denotes the Jacobian matrix of $p(Q)$, we deduce

$$\nabla \psi_f(q^N)^T q^N = p(Q^N)^T q_f^N + (q_f^N)^T Jp(Q^N) Q^N - \nabla c_f(q_f^N)^T q_f^N.$$

We also have, for all $t \neq f$,

$$\nabla_{q_t} \pi_f^*(q_{-f}^N) = Jp(Q^N)^T q_f^N,$$

which yields

$$\nabla \pi_f^*(q_{-f}^N)^T q_{-f}^N = (q_f^N)^T Jp(Q^N) Q_{-f}^N.$$

Based on the above expressions, we easily obtain the following corollary of Proposition 3, which provides a sufficient condition under which every player can increase his/her individual payoff above the Nash payoff by simply scaling the Nash equilibrium solution q^N by a factor $1 - \tau$ for $\tau > 0$ sufficiently small,

Corollary 1. *Let $c_f(q_f)$ be a convex function satisfying $c_f(0) = 0$. Assume that $c_f(q_f)$ and $p(Q)$ are continuously differentiable. If $0 \in X_f$ and*

$$\delta (-q_f^N)^T Jp(Q^N) Q_{-f}^N > \pi_f^N + (q_f^N)^T Jp(Q^N) q_f^N, \quad \forall f \in \mathcal{F},$$

then a Pareto improvement exists for the Cournot payoff function (4). In fact, $(1 - \tau)q^N$ is a Pareto improvement for all $\tau > 0$ sufficiently small.

Proof. Continuing the above derivation, we have

$$\begin{aligned} & \nabla \psi_{\delta, f}(q^N)^T q^N \\ &= \pi_f^N + c_f(q_f^N) - \nabla c_f(q_f^N)^T q_f^N + (q_f^N)^T Jp(Q^N) Q^N - (1 - \delta)(q_f^N)^T Jp(Q^N) Q_{-f}^N \\ &\leq \pi_f^N + \delta (q_f^N)^T Jp(Q^N) Q_{-f}^N + (q_f^N)^T Jp(Q^N) q_f^N < 0, \end{aligned}$$

where the first inequality follows from the convexity of c_f and the fact that $c_f(0) = 0$, and the second inequality is by assumption. Consequently, (8) holds with $q = 0$. The last assertion of the corollary is then obvious. \square

In general, the existence of a vector q satisfying the ascent condition (8) can be checked by solving the following convex optimization problem in the variable (q, τ) :

$$\begin{aligned} & \text{maximize } \tau \\ & \text{subject to } \nabla \psi_{\delta, f}(q^N)^T (q - q^N) \geq \tau, \quad \forall f \in \mathcal{F} \\ & \text{and } q \in X. \end{aligned} \tag{9}$$

Clearly, there exists no vector $q \in X$ satisfying (8) if and only if the maximum objective value of the above optimization problem is zero. If each firm f 's production set X_f is polyhedral, then (9) is a linear program. In the rest of the paper, we assume that a Pareto improvement exists.

4. The collusive optimization problem

In considering the problem of collusion, one can think of firms facing two subproblems. First, firms must identify quantity vectors that are sustainable in the sense that if they agree to produce in a certain manner then it is in the best interests of each firm to do so; in other words, they are equilibria. Second, having identified the set of sustainable quantity vectors, firms must select a particular vector from that set. This is the approach outlined and implemented in Harrington [12]. The previous section addressed the first subproblem by characterizing the set of quantity vectors supportable by a class of Nash equilibria for the infinitely repeated extension. In this section, we take on the second subproblem—selecting an element from Ω_δ . When firms are symmetric, it has been common practice to focus on the best symmetric element of Ω_δ . When firms are asymmetric, it seems natural that the selection should be asymmetric. For example, firms with more capacity (as reflected in the upper bound to X_f) or lower cost should have higher quantities. Such can be argued from a variety of perspectives; for example, the single-period Nash equilibrium has firms with lower cost producing at a higher rate so the collusive solution should also retain that property. Of course, firms will disagree over how they rank various payoff vectors. This second subproblem inherently involves bargaining among the firms—as they try to resolve their differences—which makes it natural to use a bargaining solution as a selection device.

In the axiomatic bargaining literature (e.g. [18]), a bargaining problem is defined by a set of payoff vectors, Γ , which is the set over which players bargain, and a disagreement payoff, $d \in \Gamma$, which is the payoff vector if they fail to reach an agreement. In our setting, $\Gamma = \{\pi(q) : q \in \Omega_\delta\}$ and $d = \pi^N$. For any (Γ, d) , a bargaining solution assigns a subset of Γ (preferably a singleton). The axiomatic approach to the bargaining problem is to put forth a set of axioms, which are interpreted to be desirable properties for a bargaining solution to have, and to characterize the solution that satisfies those axioms. This approach was originally laid out in Nash [17] and the Nash bargaining solution remains the most widely-accepted bargaining solution. In our setting, the Nash bargaining solution solves the following objective which we will refer to as the *Nash bargaining objective* (NBO):¹

$$\theta(q) \equiv \prod_{f \in \mathcal{F}} (\pi_f(q) - \pi_f^N), \quad q \in \Omega_\delta.$$

¹ The Nash bargaining solution is typically formulated in terms of the selection of a payoff vector:

$$\max_{\pi \in \{\pi(q) : q \in \Omega_\delta\}} \prod_{f \in \mathcal{F}} (\pi_f - \pi_f^N).$$

The difficulty with this problem is that deriving the choice set is a challenging computational task. Given that our interest ultimately lies with what firms do rather than what they earn, we have instead formulated the problem as the selection of a quantity vector:

$$\max_{q \in \Omega_\delta} \prod_{f \in \mathcal{F}} (\pi_f(q) - \pi_f^N).$$

Also, Nash originally assumed that Γ is convex; a property that need not hold for $\{\pi(q) : q \in \Omega_\delta\}$. However, Kaneko (1980) has shown that Nash's result holds as long as Γ is compact which is true in our case. Also see Herrero (1989).

The (δ) -collusive game solution is a vector that maximizes $\theta(q)$ on the set Ω_δ ; thus the collusive game problem can be stated as the following optimization problem:

$$\begin{aligned} &\text{maximize } \log \theta(q) \\ &\text{subject to } q \in \Omega_\delta. \end{aligned} \tag{10}$$

The logarithmic objective function is well defined and extended-valued (i.e., possibly equal to $-\infty$) on the feasible set (e.g., when $q = q^N$), and is finite-valued on the subset of Pareto improvements. Clearly, only the latter vectors are of interest in the solution of (10).

In addition to the extended-valued feature of the objective function, there are several computational challenges associated with (10). First and foremost is the nonconvexity of this problem in general (see the next section for more discussion). Another challenge, which endangers the superlinear convergence of computational methods for solving the problem and complicates sensitivity analysis under data perturbations, is the fact that the once but not twice differentiable implicitly defined value function $\pi_f^*(q_{-f})$ is present in the constraint set Ω_δ . Fortunately, it is not difficult to derive an equivalent formulation of the Karush-Kuhn-Tucker (KKT) system of (10) as a mixed nonlinear complementarity problem (NCP) that involves only the input functions; in particular, the resulting NCP formulation circumvents the need to evaluate the value function during computations, which is not a trivial task in realistic applications.

To derive the equivalent NCP formulation, we assume that each set X_f is finitely represented:

$$X_f \equiv \{q_f \in \mathfrak{R}^{m_f} : g_f(q_f) \leq 0\}, \tag{11}$$

where $g_f : \mathfrak{R}^{m_f} \rightarrow \mathfrak{R}^{m_f}$ is twice continuously differentiable and each component function g_{fi} for $i = 1, \dots, m_f$ is convex. In what follows, we present the KKT system of (10) without regards to prerequisite constraint qualifications (CQs); these will be formally stated when we analyze the optimization problem rigorously. (A word of caution: while the set X_f is polyhedral in many applications, the profit function $\pi_f(q)$ is always nonlinear, thereby rendering the feasible set Ω_δ non-polyhedral in all cases.) The KKT system is as follows: for all $f \in \mathcal{F}$,

$$\begin{aligned} 0 &= -\sum_{t \in \mathcal{F}} \frac{\nabla_{q_f} \pi_t(q)}{\pi_t(q) - \pi_t^N} + \sum_{i=1}^{m_f} \lambda_{fi} \nabla g_{fi}(q_f) - \mu_f \nabla_{q_f} \pi_f(q) \\ &\quad + \sum_{f \neq t \in \mathcal{F}} \mu_t \left[-\nabla_{q_f} \pi_t(q) + (1 - \delta) \nabla_{q_f} \pi_t(q_t^*(q_{-t}), q_{-t}) \right] \\ 0 &\leq \lambda_f \perp g_f(q_f) \leq 0 \\ 0 &\leq \mu_f \perp \pi_f(q) - (1 - \delta) \pi_f^*(q_{-f}) - \delta \pi_f^N \geq 0, \end{aligned}$$

where $\lambda_f \in \mathfrak{R}^{m_f}$ and $\mu_f \in \mathfrak{R}$ are the KKT multipliers of the constraints defining Ω_δ . In order to eliminate the implicit $q_f^*(q_{-f})$ and $\pi_f^*(q_{-f})$, we recall from Proposition 1 the variational characterization of the former vector, which we can state in terms of another complementarity system. Letting $\tilde{q}_f \equiv q_f^*(q_{-f})$, we can write the KKT system for firm

f 's optimization problem (2) as follows (again, we are informal here with the omission of CQs):

$$\begin{aligned} 0 &= \nabla_{q_f} \pi_f(\tilde{q}_f, q_{-f}) + \sum_{i=1}^{m_f} \tilde{\lambda}_{fi} \nabla g_{fi}(\tilde{q}_f) \\ 0 &\leq \tilde{\lambda}_f \perp g_f(\tilde{q}_f) \leq 0. \end{aligned} \quad (12)$$

Putting together the two complementarity systems, we obtain a single combined system in which the optimal response vector $q_f^*(q_{-f})$ and the value function $\pi_f^*(q_{-f})$ are both eliminated: for all $f \in \mathcal{F}$,

$$\begin{aligned} 0 &= -\sum_{t \in \mathcal{F}} \frac{\nabla_{q_f} \pi_t(q_t, q_{-t})}{\pi_t(q_t, q_{-t}) - \pi_t^N} + \sum_{i=1}^{m_f} \lambda_{fi} \nabla g_{fi}(q_f) - \mu_f \nabla_{q_f} \pi_f(q_f, q_{-f}) \\ &\quad + \sum_{f \neq t \in \mathcal{F}} \mu_t \left[-\nabla_{q_f} \pi_t(q_t, q_{-t}) + (1 - \delta) \nabla_{q_f} \pi_t(\tilde{q}_t, q_{-t}) \right] \\ 0 &= \nabla_{q_f} \pi_f(\tilde{q}_f, q_{-f}) + \sum_{i=1}^{m_f} \tilde{\lambda}_{fi} \nabla g_{fi}(\tilde{q}_f) \\ 0 &\leq \lambda_f \perp g_f(q_f) \leq 0 \\ 0 &\leq \tilde{\lambda}_f \perp g_f(\tilde{q}_f) \leq 0 \\ 0 &\leq \mu_f \perp \pi_f(q_f, q_{-f}) - (1 - \delta) \pi_f(\tilde{q}_f, q_{-f}) - \delta \pi_f^N \geq 0. \end{aligned} \quad (13)$$

The latter system, although involving the auxiliary variables \tilde{q}_f and $\tilde{\lambda}_f$, contains only the input (payoff and constraint) functions and their derivatives; there are no more implicitly defined functions.

Since an optimization subproblem (via the optimal objective value $\pi_f^*(q_{-f})$) is embedded within (10), the latter problem is a mathematical program with equilibrium constraints (MPEC) [16], or more specifically, a bilevel program. Normally, the first-order optimality conditions for such a mathematical program are rather involved; nevertheless, this is not the case with (10). The main reason is that the nondifferentiable optimal vector $q_f^*(q_{-f})$ does not appear explicitly in (10); instead, it enters through the differentiable value function $\pi_f^*(q_{-f})$. In particular, the possible non-uniqueness of the multipliers $\tilde{\lambda}_f$ in (12) is not a deterrent for the system (13) to be an equivalent formulation of the first-order optimality conditions for (10).

Convexity in the univariate case A key concern in solving the optimization problem (10) is the log-concavity of the objective function and the convexity of the feasible set. Admittedly, we have not been able to resolve this technical issue in the general case. In fact, even the analysis for the univariate case, where each player f 's decision variable is a scalar, the set X_f is the interval $[0, \text{CAP}_f]$, and the cost function $c_f(q_f) \equiv c_f q_f$ is linear, with c_f being a positive constant, is non-trivial. Instead of giving the full details of the analysis for this special case, we summarize the key result in the following proposition and give a sketch of its proof, which is somewhat involved.

Proposition 4. *In the univariate case specified above, if*

$$\pi_f^N > (1 - \delta) \beta \text{CAP}_f^2, \quad \forall \text{ firms } f \text{ such that } q_f^N < \text{CAP}_f, \quad (14)$$

then Ω_δ is a convex set and $\log \theta(q)$ is an extended-valued concave function on Ω_δ .

Sketch of the proof. The convexity proof of Ω_δ is divided into several steps. We first derive the following explicit expression:

$$\begin{aligned} & \pi_f^*(Q_{-f}) \\ &= \begin{cases} \beta \text{CAP}_f \left(\frac{\alpha - c_f}{\beta} - Q_{-f} - \text{CAP}_f \right) & \text{if } Q_{-f} \in \left[0, \frac{\alpha - c_f}{\beta} - 2\text{CAP}_f \right] \\ \beta \left(\frac{\alpha - c_f}{2\beta} - \frac{Q_{-f}}{2} \right)^2 & \text{if } Q_{-f} \in \left[\frac{\alpha - c_f}{\beta} - 2\text{CAP}_f, \frac{\alpha - c_f}{\beta} \right] \\ 0 & \text{if } Q_{-f} \in \left[\frac{\alpha - c_f}{\beta}, \infty \right), \end{cases} \end{aligned} \quad (15)$$

which shows that $\pi_f^*(Q_{-f})$ is a nonincreasing function of its argument. Since $\pi_f^*(Q_f) \geq \pi_f^*(Q_f^N)$ for $q \in \Omega_\delta$, it follows that $Q_{-f} \leq Q_{-f}^N$ for all $f \in \mathcal{F}$ and all $q \in \Omega_\delta$. This yields an equivalent representation of the set Ω_δ :

$$\Omega_\delta = \left\{ q \in \Omega_* : \pi_f(q_f, Q_{-f}) \geq (1 - \delta) \pi_f^*(Q_{-f}) + \delta \pi_f^N, \quad \forall f \in \mathcal{F} \right\},$$

where

$$\Omega_* \equiv \left\{ q \in \mathbb{R}^{|\mathcal{F}|} : 0 \leq q_f \leq \text{CAP}_f; Q_{-f} \leq \frac{\alpha - c_f}{\beta}, \quad \forall f \in \mathcal{F} \right\}$$

can be shown to contain q^N (the condition (14) is needed here). The next step is to obtain a further equivalent representation for each set:

$$\Omega_\delta^f = \left\{ q \in \Omega_* : \pi_f(q_f, Q_{-f}) \geq (1 - \delta) \pi_f^*(Q_{-f}) + \delta \pi_f^N \right\},$$

depending on whether f is a *Nash-capacitated* firm (i.e., $q_f^N = \text{CAP}_f$) or a *Nash-un-capacitated* firm (i.e., $q_f^N < \text{CAP}_f$). Specifically, one can show that for a firm of the former type, a vector $q \in \Omega_\delta^f$ if and only if $q \in \Omega_*$, $Q_{-f} \leq (\alpha - c_f)/\beta - 2\text{CAP}_f$, and

$$\begin{aligned} & \frac{\alpha - c_f}{2\beta} - \frac{Q_{-f}}{2} \geq (1 - \delta) \text{CAP}_f \\ & + \sqrt{\left[\frac{\alpha - c_f}{2\beta} - \frac{Q_{-f}}{2} - q_f \right]^2 + \frac{\delta \pi_f^N}{\beta} - \delta (1 - \delta) \text{CAP}_f^2}, \end{aligned}$$

with the square-root term being a well-defined convex function of q . For a Nash-un-capacitated firm f , we have

$$\Omega_\delta^f = \left\{ q \in \Omega_* : Q_{-f} \leq \frac{\alpha - c_f}{\beta} - 2\sqrt{\frac{\pi_f^N}{\beta}} \text{ and } q_f \in I_\delta^f(Q_{-f}) \right\},$$

where $I_\delta^f(Q_{-f})$ is the interval

$$\left[\frac{\alpha - c_f}{2\beta} - \frac{Q_{-f}}{2} - \sqrt{r_\delta(Q_{-f})}, \max \left(\text{CAP}_f, \frac{\alpha - c_f}{2\beta} - \frac{Q_{-f}}{2} + \sqrt{r_\delta(Q_{-f})} \right) \right]$$

with

$$r_\delta(Q_{-f}) \equiv$$

$$\delta \left[\left(\frac{\alpha - c_f}{2\beta} - \frac{Q_{-f}}{2} \right)^2 - \frac{\pi_f^N}{\beta} \right] + (1 - \delta) \left[\max \left(\frac{\alpha - c_f}{2\beta} - \frac{Q_{-f}}{2} - \text{CAP}_f, 0 \right) \right]^2.$$

The convexity of Ω_δ^f for a Nash-uncapacitated firm follows by showing that $\sqrt{r_\delta^f(Q_{-f})}$

is concave for Q_{-f} in the interval $\left[0, \frac{\alpha - c_f}{\beta} - 2\sqrt{\frac{\pi_f^N}{\beta}} \right]$. Since Ω_δ is the intersection

of Ω_δ^f over $f \in \mathcal{F}$, the convexity of Ω_δ follows. The log-concavity of $\theta(q)$ on Ω_δ can be established by calculating the Hessian of $-\log \theta(q)$ and showing that such a Hessian must be positive definite for $q \in \Omega_1$ such that $\pi_f(q) > \pi_f^N$ for all $f \in \mathcal{F}$. In turn, the proof of the latter assertion is based on showing that the Hessian can be written in the form $\text{Diag}(b) + (\mathbf{1}^T a)E - c\mathbf{1}^T - \mathbf{1}c^T$ for some positive vectors a, b , and c satisfying $a_i b_i > c_i^2$ for all i , where $\mathbf{1}$ and E are respectively the vector and matrix of all ones of appropriate dimensions, and that a matrix of the latter kind must be positive definite. \square

5. Bounding procedures

Returning to the multivariate setting, where each firm’s production set X_f is a compact convex subset of \mathfrak{N}^N , we focus on the affine payoff function (5). Without loss of generality, we may take β to be the vector of all ones by a simple scaling of the data and the variables. Noting the identity

$$\pi_f(q) = \left\| \frac{\alpha}{2} - \frac{Q_{-f}}{2} \right\|^2 - \left\| \frac{\alpha}{2} - \frac{Q_{-f}}{2} - q_f \right\|^2 - c_f(q_f), \tag{16}$$

we see that $q \in \Omega_\delta$ if and only if $q \in X$ and for all $f \in \mathcal{F}$, (since $Q + q_f = Q_{-f} + 2q_f$)

$$\| \alpha - Q_{-f} \|^2 \geq \| \alpha - Q - q_f \|^2 + 4 \left[c_f(q_f) + (1 - \delta) \pi_f^*(q_{-f}) + \delta \pi_f^N \right]. \tag{17}$$

Based on this “dc” (difference of convex) representation of the set Ω_δ , we develop iterative upper and lower bounding procedures for dealing with the optimization problem (10). It should be noted that while there is an extensive literature on dc programming (see, e.g., the recent review article [14]), due to the special structure of the problem (10), we choose to develop the bounding procedures from first principles that are more akin to the problem on hand. As an example of an alternative iterative bounding procedure, we mention the “DCA” described in [24], where duality plays a key role. For the latter algorithm to be directly applicable, it is imperative that the constraints be convex. Since we cannot guarantee the convexity of (17), the direct application of the DCA is in jeopardy.

5.1. Upper bounding

The boundedness of the set X induces an upper bound, which we denote κ_f , on the quantity $\|\alpha - Q_{-f}\|^2$ for each $f \in \mathcal{F}$. Alternatively, if X_f is contained in the non-negative orthant (which is a typical situation in applications) and if the cost function c_f is nonnegative, then the constraint (17) itself trivially induces an upper bound on each $\|q_f\|$, and thus on $\|\alpha - Q_{-f}\|^2$ for all $f \in \mathcal{F}$. In what follows, we take the bound κ_f to be a readily available quantity.

Consider the following set which involves the auxiliary variables ξ_f for $f \in \mathcal{F}$:

$$\begin{aligned} \Gamma_\delta &\equiv \{ (q, \xi) \in X \times \mathbb{R}^{|\mathcal{F}|} : \text{for all } f \in \mathcal{F}, \quad \xi_f \leq \kappa_f \\ &\quad \xi_f \geq \|\alpha - Q - q_f\|^2 + 4[c_f(q_f) + (1 - \delta)\pi_f^*(Q_{-f}) + \delta\pi_f^N] \\ &\quad \xi_f \geq \|\alpha - Q_{-f}\|^2 \\ &\quad \xi_f \geq \|\alpha - Q - q_f\|^2 + 4[c_f(q_f) + \pi_f^N] \}. \end{aligned}$$

It is clear that Γ_δ is a “relaxation” of Ω_δ in the sense that if $q \in \Omega_\delta$, then (q, ξ) , where $\xi_f \equiv \|\alpha - Q_{-f}\|^2$ for all f , belongs to Γ_δ . Moreover, Γ_δ is clearly a convex set. To recover the set Ω_δ from Γ_δ , we consider the penalization of the next-to-last inequality in Γ_δ . Specifically, for each scalar $\varsigma > 0$, we consider the optimization problem in the variable (q, ξ) :

$$\begin{aligned} &\text{maximize } \sum_{f \in \mathcal{F}} \left\{ \log \left(\frac{1}{4} \left[\xi_f - \|\alpha - Q - q_f\|^2 - 4(\pi_f^N + c_f(q_f)) \right] \right) \right. \\ &\quad \left. \varsigma \sum_{f \in \mathcal{F}} \left(\|\alpha - Q_{-f}\|^2 - \xi_f \right) \right\} \\ &\text{subject to } q \in X \end{aligned} \tag{18}$$

$$\begin{aligned} \text{and } \quad &\forall f \in \mathcal{F}, \quad \xi_f \leq \kappa_f \\ &\xi_f \geq \|\alpha - Q - q_f\|^2 + 4 \left[(1 - \delta)\pi_f^*(Q_{-f}) + \delta\pi_f^N + c_f(q_f) \right] \\ &\xi_f \geq \|\alpha - Q - q_f\|^2 + 4 \left[\pi_f^N + c_f(q_f) \right] \\ &\xi_f \geq \|\alpha - Q_{-f}\|^2. \end{aligned}$$

By the next-to-last constraint, the logarithmic term in the objective function is a well-defined extended-valued concave function, with value equal to $-\infty$ when equality holds in the next-to-last constraint; the concavity is easy to verify. The penalty term $\varsigma \sum_{f \in \mathcal{F}} (\|\alpha - Q_{-f}\|^2 - \xi_f)$ in the objective function, together with the final constraint,

is intended to force $\xi_f - \|\alpha - Q_{-f}\|^2$ to zero, hence the log part of the objective function to $\log \theta(q)$ as $\varsigma \rightarrow \infty$. Note that the objective of (18) is a dc function because of the term $\|\alpha - Q_{-f}\|^2$. The next result summarizes several key properties of the above optimization problem whose maximum objective value we denote ω_ς .

Proposition 5. Let X_f be a nonempty compact convex subset of \mathbb{R}^N containing the origin, and let c_f be a nonnegative function. The following statements are valid.

- (a) For every $\varsigma > 0$, $\omega_\varsigma \geq \max_{q \in \Omega_\delta} \log \theta(q) \equiv \omega_\infty$;
 (b) $\varsigma_1 > \varsigma_2 > 0$ implies $\omega_{\varsigma_2} \geq \omega_{\varsigma_1}$, with equality holding if and only if $\omega_{\varsigma_2} = \omega_\infty$;
 (c) $\lim_{\varsigma \rightarrow \infty} \omega_\varsigma = \omega_\infty$.

Proof. The proposition is a type of exact penalty function result for a constrained optimization problem. For completeness, we give the details of the proof of parts (b) and (c); part (a) does not require a proof. Let $\varsigma_1 > \varsigma_2 > 0$ and let (q^ν, ξ^ν) denote an optimal solution of (18) corresponding to ς_ν for $\nu = 1, 2$. We have

$$\begin{aligned} \omega_{\varsigma_2} &\geq \sum_{f \in \mathcal{F}} \left\{ \log \left(\frac{1}{4} \left[\xi_f^1 - \|\alpha - Q^1 - q_f^1\|^2 - 4(\pi_f^N + c_f(q_f^1)) \right] \right) \right. \\ &\quad \left. + \varsigma_2 \sum_{f \in \mathcal{F}} \left(\|\alpha - Q_{-f}^1\|^2 - \xi_f^1 \right) \right\} \\ &\geq \sum_{f \in \mathcal{F}} \left\{ \log \left(\frac{1}{4} \left[\xi_f^1 - \|\alpha - Q^1 - q_f^1\|^2 - 4(\pi_f^N + c_f(q_f^1)) \right] \right) \right. \\ &\quad \left. + \varsigma_1 \sum_{f \in \mathcal{F}} \left(\|\alpha - Q_{-f}^1\|^2 - \xi_f^1 \right) \right\} \\ &= \omega_{\varsigma_1} \end{aligned}$$

If $\omega_{\varsigma_2} = \omega_{\varsigma_1}$, then $\|\alpha - Q_{-f}^1\|^2 = \xi_f^1$ for all $f \in \mathcal{F}$. Hence q^1 belongs to Ω_δ . Consequently, $\omega_\infty = \omega_{\varsigma_1}$; and so $\omega_{\varsigma_2} = \omega_\infty$. The converse is obvious.

Let $\{\varsigma_k\}$ be any sequence of scalars tending to ∞ and let (q^k, ξ^k) be a corresponding sequence of optimal solutions of (18). The latter sequence of optimal solutions must be bounded. If (q^∞, ξ^∞) is any accumulation point of this sequence, then we must have $\|\alpha - Q_{-f}^\infty\|^2 = \xi_f^\infty$ for all $f \in \mathcal{F}$ because the sequence of optimal objective values $\{\omega_{\varsigma_k}\}$ is nonincreasing and bounded below by ω_∞ , and thus converges. This shows that q^∞ belongs to Ω_δ . \square

Ideally, we would like to have a convergent upper bounding procedure that solves convex optimization subproblems. In addition to a dc programming approach, one way to achieve this is via a global optimization method, such as that of convex extension and enveloping [25, 26], possibly embedded within a “branch-and-cut scheme”. Detailed exploration of such a scheme is beyond the scope of this paper and is under development in the Ph.D. thesis of Liu.

5.2. Lower bounding

Instead of replacing the left-hand side of (17) by ξ_f and then penalizing the resulting inequality $\xi_f \geq \|\alpha - Q_{-f}\|^2$, we use a first-order approximation of the term $\|\alpha - Q_{-f}\|^2$ for the derivation of lower bounds for the collusive game optimization problem (10).

We recall the blanket assumption that a Pareto improvement q^0 exists. By definition, q^0 belongs to the set Ω_δ and satisfies

$$\pi_f^*(Q_{-f}^0) > \pi_f^N, \quad \forall f \in \mathcal{F}. \quad (19)$$

If no such vector q^0 exists, then the maximization problem (10) is not interesting at all because its objective function is identically equal to zero on its feasible set. Noting the equality

$$\|\alpha - Q_{-f}\|^2 = \|\alpha - Q_{-f}^0\|^2 + 2(Q_{-f}^0 - \alpha)^T(Q_{-f} - Q_{-f}^0) + \|Q_{-f} - Q_{-f}^0\|^2,$$

we see that the set

$$\begin{aligned} \Omega_\delta(q^0) &\equiv \{q \in X : \|\alpha - Q_{-f}\|^2 + 2(Q_{-f}^0 - \alpha)^T(Q_{-f} - Q_{-f}^0) \\ &\geq \|\alpha - Q - q_f\|^2 + 4[c_f(q_f) + (1 - \delta)\pi_f^*(Q_{-f}) + \delta\pi_f^N] \quad \forall f\} \end{aligned}$$

is a nonempty subset of Ω_δ , nonempty because $\Omega_\delta(q^0)$ contains q^0 . Moreover, the set $\Omega_\delta(q^0)$ is convex because the left-hand side of the inequality in $\Omega_\delta(q^0)$ is a linear function of Q_{-f} and the right-hand side is a convex function of (q_f, Q_{-f}) . In essence, $\Omega_\delta(q^0)$ is obtained by “semi-linearizing” the quadratic function $\pi_f(q_f, Q_{-f})$, whereby the first term in the right-hand side of (16) is linearized at Q^0 and the second term is left intact.

The next result shows that the inequality (19) continues to hold for all elements of the set $\Omega_\delta(q^0)$.

Proposition 6. *For all $q \in \Omega_\delta(q^0)$, $\pi_f^*(Q_{-f}) > \pi_f^N$ for all $f \in \mathcal{F}$.*

Proof. Let $q \in \Omega_\delta(q^0)$. We know that $\pi_f^*(Q_{-f}) \geq \pi_f^N$ for all $f \in \mathcal{F}$. Assume for the sake of contradiction that $\pi_f^*(Q_{-f}) = \pi_f^N$ for some $f \in \mathcal{F}$. We then have

$$\pi_f^N \geq \pi_f(q) \geq \frac{1}{4} \|Q_{-f} - Q_{-f}^0\|^2 + (1 - \delta)\pi_f^*(Q_{-f}) + \delta\pi_f^N,$$

which implies $Q_{-f} = Q_{-f}^0$. This contradicts (19). \square

Similar to the restricted set $\Omega_\delta(q^0)$, we consider, for a given vector q^0 satisfying (19), a lower objective function:

$$\begin{aligned} \chi(q, q^0) &\equiv \prod_{f \in \mathcal{F}} \frac{1}{4} \left[\|\alpha - Q_{-f}^0\|^2 + 2(Q_{-f}^0 - \alpha)^T(Q_{-f} - Q_{-f}^0) \right. \\ &\quad \left. - \|\alpha - Q - q\|^2 - 4(\pi_f^N + c_f(q_f)) \right], \end{aligned}$$

which has the property that

$$\chi(q, q) = \theta(q) \geq \chi(q, q^0), \quad \forall q \in \Omega_\delta(q^0). \quad (20)$$

Taking logarithm of the above lower objective function, we consider the restricted collusive optimization problem:

$$\begin{aligned} & \text{maximize } \log \chi(q, q^0) \\ & \text{subject to } q \in \Omega_\delta(q^0). \end{aligned} \quad (21)$$

Defining

$$\varphi_f(q, q^0) \equiv \frac{1}{4} \left[\|\alpha - Q_{-f}^0\|^2 + 2(Q_{-f}^0 - \alpha)^T (Q_{-f} - Q_{-f}^0) - \|\alpha - Q - q\|^2 - 4(\pi_f^N + c_f(q_f)) \right],$$

we see that $\varphi_f(\cdot, q^0)$ is a concave in the first variable, and that

$$\chi(q, q^0) = \prod_{f \in \mathcal{F}} \varphi_f(q, q^0).$$

Moreover, $\varphi_f(q^0, q^0) = \pi_f(q^0) - \pi_f^N$. The following result summarizes some basic properties of the optimization problem (20) in the case of linear cost function $c_f(q_f) \equiv c_f^T q_f$, where c_f is a given vector with components c_{fi} .

Proposition 7. *Suppose that X_f is a closed convex subset of \mathbb{R}^N and that $c_{fi} \geq 0$ for all f and i . Let $\delta \in (0, 1)$ and $q^0 \in \Omega_\delta$ satisfying (19) be given. The following three statements hold.*

- (a) *The function $\chi(\cdot, q^0)$ is positive and $\log \chi(\cdot, q^0)$ is strictly concave on $\Omega_\delta(q^0)$; thus, (21) is a well-defined concave maximization problem.*
- (b) *Problem (21) has a unique optimal solution, say q^1 , which satisfies*

$$\max_{q \in \Omega_\delta} \theta(q) \geq \theta(q^1) \geq \max_{q \in \Omega_\delta(q^0)} \chi(q, q^0) \geq \theta(q^0). \quad (22)$$

Moreover, $\theta(q^1) = \theta(q^0)$ if and only if $q^1 = q^0$.

- (c) *If either [X_f is finitely represented and the Mangasarian-Fromovitz constraint qualification (MFCQ) holds for the set Ω_δ at q^0], or [X_f is polyhedral], then $\theta(q^1) = \theta(q^0)$ if and only if q^0 is a KKT point of the original collusive game optimization problem (10).*

Proof. For q in $\Omega_\delta(q^0)$, we have

$$\varphi_f(q, q^0) \geq (1 - \delta) (\pi_f^*(Q_{-f}) - \pi_f^N).$$

Therefore, by Proposition 6, it follows that $\chi(q, q^0)$, being the product of $\varphi_f(q, q^0)$ for all $f \in \mathcal{F}$, is positive on $\Omega_\delta(q^0)$. Since

$$\log \chi(q, q^0) = \sum_{f \in \mathcal{F}} \log \varphi_f(q, q^0).$$

it follows that $\log \chi(\cdot, q^0)$ is concave on $\Omega_\delta(q^0)$. To show the strict concavity, let $\tau \in (0, 1)$ and let q and q' be two distinct vectors in $\Omega_\delta(q^0)$. For the sake of contradiction, assume that

$$\log \chi(\tau q + (1 - \tau)q', q^0) = \tau \log \chi(q, q^0) + (1 - \tau) \log \chi(q', q^0).$$

Then we must have, for all $f \in \mathcal{F}$,

$$\varphi_f(\tau q + (1 - \tau)q', q^0) = \tau \varphi_f(q, q^0) + (1 - \tau) \varphi_f(q', q^0),$$

or equivalently,

$$\begin{aligned} & \| \alpha - c_f - \tau Q - (1 - \tau) Q' - \tau q_f - (1 - \tau) q'_f \|^2 = \\ & \tau \| \alpha - c_f - Q - q_f \|^2 + (1 - \tau) \| \alpha - c_f - Q' - q'_f \|^2. \end{aligned}$$

By the strict convexity of the Euclidean norm, it follows that $Q + q_f = Q' + q'_f$ for all $f \in \mathcal{F}$. Summing up these equalities over all $f \in \mathcal{F}$ yields $Q = Q'$, which in turn implies $q = q'$. Consequently, the strict concavity of $\log \chi(\cdot, q^0)$ follows. This establishes (a).

The existence and uniqueness of q^1 do not require a proof. The first inequality in (22) is obvious because $\Omega_\delta(q^0)$ is a subset of Ω_δ ; the second inequality is due to (20); and the third inequality is because $\Omega_\delta(q^0)$ contains q^0 . If $\theta(q^1) = \theta(q^0)$, then we must have $\chi(q^1, q^0) = \chi(q^0, q^0)$, which shows that q^0 is an optimal solution of (21). By the uniqueness of q^1 , it follows that $q^1 = q^0$. Hence (b) holds.

Under the assumptions of part (c), the KKT conditions are necessary optimality conditions for (21). Hence, for X_f represented by (11), there exist multipliers $\lambda_f \in \Re^{m_f}$ and $\eta_f \in \Re$ such that, for all $f \in \mathcal{F}$,

$$\begin{aligned} 0 &= \sum_{f \neq t \in \mathcal{F}} \frac{Q^1 + q_t^1 - Q_{-t}^0}{2 \varphi_t(q^1, q^0)} + \frac{Q^1 + q_f^1 - \alpha + \nabla c_f(q_f^1)}{\varphi_f(q^1, q^0)} + \sum_{i=1}^{m_f} \lambda_{fi} \nabla g_{fi}(q_f^1) \\ &+ 4 \eta_f \left[q_f^1 + Q^1 - \alpha + \nabla c_f(q_f^1) \right] \\ &+ \sum_{f \neq t \in \mathcal{F}} 2 \eta_t \left[(2q_t^1 + Q_{-t}^1 - Q_{-t}^0) - 4(1 - \delta) q_t^*(Q_{-t}^1) \right] \\ 0 &\leq \lambda_f \perp g_f(q_f^1) \leq 0 \\ 0 &\leq \eta_f \perp \| \alpha - Q_{-f}^0 \|^2 + 2(Q_{-f}^0 - \alpha)^T (Q_{-f}^1 - Q_{-f}^0) \\ &- \| \alpha - Q^1 - q_f^1 \|^2 - 4[c_f(q_f^1) + (1 - \delta) \pi_f^*(Q_{-f}^1) + \delta \pi_f^N] \geq 0. \end{aligned}$$

If $\theta(q^1) = \theta(q^0)$, then $q^1 = q^0$. Hence the above KKT system reduces to

$$\begin{aligned}
0 &= \sum_{i \in \mathcal{F}} \frac{q_i^0}{\pi_i(q^0) - \pi^N} + \frac{Q^0 - \alpha + \nabla c_f(q_f^0)}{\pi_f(q^0) - \pi^N} + \sum_{i=1}^{m_f} \lambda_{fi} \nabla g_{fi}(q_f^0) \\
&\quad + 4 \eta_f \left[q_f^0 + Q^0 - \alpha + \nabla c_f(q_f^0) \right] + 4 \sum_{f \neq i \in \mathcal{F}} \eta_i \left[q_i^1 - (1 - \delta) q_i^*(Q_{-i}^0) \right] \\
0 &\leq \lambda_f \perp g_f(q_f^0) \leq 0 \\
0 &\leq \eta_f \perp \pi_f(q^0) - (1 - \delta) \pi_f^*(Q_{-f}^0) - \delta \pi_f^N \geq 0.
\end{aligned}$$

With a moment's verification, the reader can easily see that the above is precisely the KKT system of (10) at q^0 . Conversely, if q^0 satisfies the latter KKT system, then $q^1 \equiv q^0$ satisfies the former KKT system. Since (21) is a concave maximization problem, it follows that q^0 is an optimal solution to it. \square

Based on Proposition 7, we propose an iterative algorithm for computing a KKT point of (10). The algorithm requires an initial vector $q^0 \in \Omega_\delta$ satisfying (19). Such a vector can be computed either by first solving the convex program (9) followed by an Armijo-type line search as instructed by Proposition 3, or by employing Corollary 1 if the latter is directly applicable.

An Iterative Lower Bounding Algorithm.

Step 0 Assume that a vector $q^0 \in \Omega_\delta$ satisfying (19) is given. Let $k = 0$.

Step 1 Solve the concave maximization subproblem:

$$\begin{aligned}
&\text{maximize } \log \chi(q, q^k) \\
&\text{subject to } q \in \Omega_\delta(q^k),
\end{aligned} \tag{23}$$

and let q^{k+1} be the unique optimal solution.

Step 2 If $\theta(q^{k+1}) - \theta(q^k)$ is less than a prescribed tolerance, stop. Otherwise, let $k \leftarrow k + 1$ and return to Step 1.

Consider an infinite sequence $\{q^k\}$ generated by the algorithm. This sequence is contained in the bounded set Ω_δ . It therefore has at least one accumulation point. The sequence of positive scalars $\{\theta(q^k)\}$ is strictly increasing and bounded above; it therefore converges. Consequently, we have

$$\lim_{k \rightarrow \infty} \theta(q^{k+1}) = \lim_{k \rightarrow \infty} \chi(q^{k+1}, q^k).$$

Since this common limit is bounded below by the positive scalar $\theta(q^0)$, it follows that

$$\lim_{k \rightarrow \infty} \|Q_{-f}^{k+1} - Q_{-f}^k\| = 0, \quad \forall f \in \mathcal{F}. \tag{24}$$

Based on the above limit, we can establish the desired KKT property of an accumulation point of the sequence $\{q^k\}$.

Proposition 8. *Suppose that X_f is a closed convex, finitely represented subset of \mathfrak{R}^N and that $c_f(q_f) = c_f^T q_f$, where each constant c_{f_i} is nonnegative. Let $\delta \in (0, 1)$ and $q^0 \in \Omega_\delta$ satisfy (19). The following two statements are valid.*

- (a) *Every accumulation point of the sequence $\{q^k\}$ belongs to Ω_δ .*
 (b) *If q^∞ is the limit of a convergent subsequence $\{q^{k+1} : k \in \mathcal{K}\}$ for some infinite subset \mathcal{K} such that q^∞ satisfies the MFCQ for the set Ω_δ , then q^∞ is a KKT point of (10).*

Proof. It suffices to prove (b). Since q^∞ belongs to Ω_δ , it follows that $\pi_f(q^\infty) \geq \pi_f^N$ for all $f \in \mathcal{F}$. Since $\theta(q^k) \geq \theta(q^0) > 0$ for all k , we must have $\pi_f(q^\infty) > \pi_f^N$ for all $f \in \mathcal{F}$. Since q^{k+1} satisfies

$$\begin{aligned} \pi_f(q^{k+1}) &- \frac{1}{4} \|Q_{-f}^{k+1} - Q_{-f}^k\|^2 \\ &= \frac{1}{4} \left[\|\alpha - Q_{-f}^k\|^2 + 2(Q_{-f}^k - \alpha)^T (Q_{-f}^{k+1} - Q_{-f}^k) - \|\alpha - Q_{-f}^{k+1} - q_f^{k+1}\|^2 \right] \\ &\quad - c_f(q_f^{k+1}) \\ &\geq (1 - \delta) \pi_f^*(Q_{-f}^{k+1}) + \delta \pi_f^N, \end{aligned}$$

and since the MFCQ is invariant under small function perturbations, it follows from (24) that the KKT conditions for (21) must necessarily be satisfied by its optimal solution q^{k+1} . Therefore, there exist multipliers $\lambda_f^{k+1} \in \mathfrak{R}^{m_f}$ and $\eta_f^{k+1} \in \mathfrak{R}$ such that, for all $f \in \mathcal{F}$,

$$\begin{aligned} 0 &= \sum_{f \neq t \in \mathcal{F}} \frac{Q^{k+1} + q_t^{k+1} - Q_{-t}^k}{2 \varphi_t(q^{k+1}, q^k)} + \frac{Q^{k+1} + q_f^{k+1} - \alpha + \nabla c_f(q_f^{k+1})}{\varphi_f(q^{k+1}, q^k)} \\ &\quad + \sum_{i=1}^{m_f} \lambda_{f_i}^{k+1} \nabla g_{f_i}(q_f^{k+1}) + 4 \eta_f^{k+1} \left[q_f^{k+1} + Q^{k+1} - \alpha + \nabla c_f(q_f^{k+1}) \right] \\ &\quad + \sum_{f \neq t \in \mathcal{F}} \eta_t^{k+1} \left[2(2q_t^{k+1} + Q_{-t}^{k+1} - Q_{-t}^k) - 4(1 - \delta) q_t^*(Q_{-t}^{k+1}) \right] \\ 0 &\leq \lambda_f^{k+1} \perp g_f(q_f^{k+1}) \leq 0 \\ 0 &\leq \eta_f^{k+1} \perp \|\alpha - Q_{-f}^k\|^2 + 2(Q_{-f}^k - \alpha)^T (Q_{-f}^{k+1} - Q_{-f}^k) \\ &\quad - \|\alpha - Q^{k+1} - q_f^{k+1}\|^2 - 4 \left[c_f(q_f^{k+1}) + (1 - \delta) \pi_f^*(Q_{-f}^{k+1}) + \delta \pi_f^N \right] \geq 0. \end{aligned}$$

From (24), we deduce

$$\lim_{k(\in \mathcal{K}) \rightarrow \infty} \varphi_f(q^{k+1}, q^k) = \pi_f(q^\infty) - \pi_f^N > 0.$$

The sequences of multipliers $\{\lambda_f^{k+1} : k \in \mathcal{K}\}$ and $\{\eta_f^{k+1} : k \in \mathcal{K}\}$ are bounded, by the MFCQ. Without loss of generality, assume that these two sequences of multipliers

converge to λ_f^∞ and η_f^∞ , respectively. To show that q^∞ is a KKT point of (10), we note that

$$\pi_f(q^\infty) = \lim_{k \in \mathcal{K} \rightarrow \infty} \left\{ \frac{1}{4} \left[\|\alpha - Q_{-f}^k\|^2 + 2(Q_{-f}^k - \alpha)^T (Q_{-f}^{k+1} - Q_{-f}^k) - \|\alpha - Q^{k+1} - q_f^{k+1}\|^2 \right] - c_f(q_f^{k+1}) \right\}$$

and

$$\lim_{k \in \mathcal{K} \rightarrow \infty} q_t^*(Q_{-t}^{k+1}) = q_t^*(Q_{-t}^\infty),$$

where the limit is due to the continuity of the optimal response function. Passing to the limit $k \in \mathcal{K} \rightarrow \infty$ in the above KKT conditions for (21), we deduce

$$0 = \sum_{t \in \mathcal{F}} \frac{q_f^\infty}{\pi_f(q^\infty) - \pi_f^N} + \frac{Q^\infty - \alpha + \nabla c_f(q_f^\infty)}{\pi_f(q^\infty) - \pi_f^N} + \sum_{i=1}^{m_f} \lambda_{fi}^\infty \nabla g_{fi}(q_f^\infty) + 4\eta_f^\infty \left[q_f^\infty + Q^\infty - \alpha + \nabla c_f(q_f^\infty) \right] + 4 \sum_{f \neq t \in \mathcal{F}} \eta_t^\infty \left[q_t^\infty - (1 - \delta) q_t^*(Q_{-t}^\infty) \right]$$

$$0 \leq \lambda_f^\infty \perp g_f(q^\infty) \leq 0$$

$$0 \leq \eta_f^\infty \perp \pi_f(q^\infty) - (1 - \delta) \pi_f^*(Q_{-f}^\infty) - \delta \pi_f^N \geq 0.$$

It can be verified that the latter system is precisely the set of KKT conditions for (10). □

6. Computational results

In this section we report computational results with the numerical solution of some test problems to illustrate the differences between the Nash quantities and prices and those obtained by collusion and to demonstrate the effectiveness of the two bounding procedures. We ran three sets of experiments. The first set pertains to univariate firms where Proposition 4 is applicable; the second and third set of runs pertain to firms with multivariate decision variables and are aimed at testing the effectiveness of the bounding schemes in comparison with some publicly available nonlinear program (NLP) solvers that are available on the NEOS website; see below for details.

6.1. A univariate problem

Demonstrably a concave maximization problem, the first test problem has 6 univariate firms. The price function is given by

$$p(Q) \equiv 40 - 0.08 Q;$$

the other data for the problem are summarized in Table 1. We wrote a simple AMPL program [8] to compute the Nash quantities q_f^N for the firms and submitted it to the

Table 1. A 6-firm problem with $\delta = 0.6$

Firm	CAP_f	c_f	q_f^N	π_f^N	capacitated?	q_f^{opt}	π_f^{opt}	$\pi_f^*(Q_{-f})$	incentive compat. binding?
A	60	15	60	516	Yes	44.795	604.030	736.10	Yes
B	20	15	20	172	Yes	18.329	247.158	267.00	No
C	55	20	45	162	No	27.818	236.014	347.00	Yes
D	48	20	45	162	No	26.822	227.570	325.90	Yes
E	25	20	25	90	Yes	16.182	137.190	194.90	No
F	10	15	10	86	Yes	10	134.843	134.80	No

NEOS server at Argonne National Laboratory (<http://www-neos.mcs.anl.gov/neos/>) for solution by the PATH solver, which was jointly written and maintained by Michael Ferris at the University of Wisconsin, Madison and Todd Munson at Argonne. We next submitted an AMPL code to solve the collusive game optimization problem with $\delta = 0.6$ and obtained the same solution using three NLP solvers: MINOS, SNOPT and LOQO. The optimal value function $\pi_f^*(Q_{-f})$ is coded according to the explicit expression (15). Although this function is only once but not twice continuously differentiable, the NLP solvers have no difficulty dealing with it. The results (rounded to 3 decimal places) are summarized in Table 1. For this example, firm C and D are Nash-uncapacitated. More importantly, Proposition 4 is applicable to this example so that the solution reported in the table is indeed globally optimal. For this problem, all firms' (except firm F) collusive production quantities q_f^{opt} are less than their respective Nash quantities, and as expected all collusive profits π_f^{opt} exceed the respective single-period Nash profits.

6.2. A multivariate problem with decoupled capacities

The next test problem has five firms constrained only by their production capacities in each of three regions. The production costs and capacities of each firm and the price and demand intercepts at each node are all given in Table 2. Thus $\alpha_i \equiv P_i^0$ and $\beta \equiv P_i^0/Q_i^0$ for $i = 1, 2, 3$. The Nash equilibrium solution for this problem is obtained by PATH and is reported in the column labelled Nash in Table 3, which also contains the results from the lower bounding and upper bounding procedures. Some details on the implementation of the latter procedures are as follows. The starting iterate for the lower bounding scheme is obtained by lowering the Nash equilibrium solution slightly, resulting in

$$q_1^0 = \begin{pmatrix} 55 \\ 18 \\ 45 \end{pmatrix}, q_2^0 = \begin{pmatrix} 46 \\ 30 \\ 57 \end{pmatrix}, q_3^0 = \begin{pmatrix} 96 \\ 75 \\ 35 \end{pmatrix}, q_4^0 = \begin{pmatrix} 21 \\ 7 \\ 16 \end{pmatrix}, \text{ and } q_5^0 = \begin{pmatrix} 28 \\ 53 \\ 97 \end{pmatrix}. \tag{25}$$

Each subproblem (23) is solved by PATH applied to its KKT conditions, utilizing a set of auxiliary variables to handle the profit function $\pi_f^*(Q_{-f}^{k+1})$ as described in Section 4.

Table 2. Data for test problem 2

Node	(Firm, node)	c_{fi}	CAP_{fi}	P_i^0	Q_i^0
1				40	500
2				35	400
3				32	640
	(1,1)	15	300		
	(1,2)	15	20		
	(1,3)	15	50		
	(2,1)	16	50		
	(2,2)	16	400		
	(2,3)	16	100		
	(3,1)	12	200		
	(3,2)	12	100		
	(3,3)	12	40		
	(4,1)	18	300		
	(4,2)	18	20		
	(4,3)	18	150		
	(5,1)	14	30		
	(5,2)	14	100		
	(5,3)	14	400		

The results in the lower-bound column in Table 3 are obtained after 28 iterations, with the relative difference of the Nash bargaining objective at the last two iterations prior to termination being:

$$\frac{\theta(q^{k+1}) - \theta(q^k)}{\theta(q^{k+1})} \approx 3.3582 \times 10^{-6}.$$

Since the firms' constraints are only regional production capacities, the optimal profit functions $\pi_f^*(Q_{-f})$ can again be explicitly represented. However, since there are three regions involved, each firm f has three decision variables q_{fi} for $i = 1, 2, 3$; more importantly, the overall collusive game optimization problem (10) is not shown to be a concave maximization problem. Nevertheless, most NLP solvers on NEOS were able to solve the problem; in particular, SNOPT produces a solution that is very close to that from the bounding procedures. The column labelled Upper Bound in Table 3 is obtained from solving the penalized problem (18), modified in the following way. Using the vector q^0 in (25), we compute $\pi_f^*(q_{-f}^0)$ for all $f \in \mathcal{F}$. We then fix all these optimal profits and solve the resulting problem (18) by the NLP solver, SNOPT6.1, using $\varsigma = 10^{10}$, obtaining the solution presented in the upper-bound column in Table 3. While the sum of the penalty terms after scaling the β_i to be unity, $\sum_{f \in \mathcal{F}} \left(\|\alpha - Q_{-f}\|^2 - \xi_f \right)$, is -0.09715

at termination, the computed solution in the modified upper bound problem nevertheless coincides with that obtained from a direct solution of the problem (10) by an NLP solver on NEOS.

A major motivation in solving this test problem in different ways is to assess the quality of the solutions produced by the approaches. This assessment is useful because of the likely nonconvexity of the collusive optimization problem (10) being solved. With the three objective function values being very close to each other (cf. the last row

Table 3. Numerical results on test problem 2 with $\delta = 0.8$

	Variable	Nash	Lower Bound	Upper Bound	NLP
Generation	q_{11}	59	55.7603	55.7113	55.7113
	q_{12}	20	11.6396	11.6915	11.6915
	q_{13}	50	2.6264	2.6231	2.62312
	total	129	70.0263	70.0259	70.0259
	q_{21}	46.5	50	50	50
	q_{22}	30.2857	11.3264	11.2944	11.2944
	q_{23}	57.5	0	0	0
	total	134.2857	61.3263	61.2944	61.2944
	q_{31}	96.5	15.1505	15.1600	15.1600
	q_{32}	76	100	100	100
	q_{33}	40	40	40	40
	total	212.5	155.1505	155.1600	155.1600
	q_{41}	21.5	29.3681	29.3557	29.3557
	q_{42}	7.42875	0	0	0
	q_{43}	17.5	0	0	0
total	46.42875	29.3681	29.3557	29.3557	
q_{51}	30	0	0	0	
q_{52}	53.1429	0	0	0	
q_{53}	97.5	138.9204	138.915	138.915	
total	180.6429	138.9204	138.915	138.915	
Profits	firm 1	545.23	852.0030	852.0331	852.0331
	firm 2	418.5495	692.2188	692.1436	692.1436
	firm 3	1525.3798	1903.0242	1903.0818	1903.0818
	firm 4	57.1211	293.0262	293.0239	293.0239
	firm 5	894.0269	1239.5394	1239.5516	1239.5516
NBO	$\theta(q)$	0	2.584208E12	2.584210E12	2.584210E12

in Table 3), we are fairly confident that the obtained solution, which is not shown to be globally optimal, is at least reasonable. This test also demonstrates the effectiveness of the two bounding procedures.

6.3. A multivariate problem with coupling capacities

Our third test problem involves the following coupling capacity constraint for each firm:

$$\sum_{i \in \mathcal{N}} q_{fi} \leq \text{CAP}_f. \tag{26}$$

The data are the same as those in test problem 2; in addition, $\text{CAP}_1 = 200$, $\text{CAP}_2 = 50$, $\text{CAP}_3 = 200$, $\text{CAP}_4 = 150$, and $\text{CAP}_5 = 110$. The starting iterate for the lower bounding scheme is obtained in the same way as before and is given as follows:

$$q_1^0 = \begin{pmatrix} 67 \\ 51 \\ 80 \end{pmatrix}, q_2^0 = \begin{pmatrix} 24 \\ 12 \\ 12 \end{pmatrix}, q_3^0 = \begin{pmatrix} 67 \\ 51 \\ 80 \end{pmatrix}, q_4^0 = \begin{pmatrix} 35 \\ 22 \\ 30 \end{pmatrix}, \text{ and } q_5^0 = \begin{pmatrix} 41 \\ 28 \\ 30 \end{pmatrix}.$$

In addition to the lower bounding procedure, we also ran the upper bounding procedure in the same way as in test problem 2. We have verified that the solution computed by the latter procedure is feasible to (10); nevertheless, since the modified problem is not the true upper bound problem (18), we cannot conclude positively that the latter solution is globally optimal to (10). See Table 4.

Finally, for comparison purposes, we use an NLP solver to directly solve an MPEC formulation of the collusive game problem (10). Due to the aggregate capacities (26), the profit function $\pi_f^*(Q_{-f})$ no longer has an explicit representation. In order to deal with this issue, we first observe that the problem (10) is equivalent to the following MPEC:

$$\begin{aligned} & \text{maximize} \quad \prod_{f \in \mathcal{F}} \left(\pi_f(q) - \pi_f^N \right) \\ & \text{subject to} \quad \pi_f(q) \geq (1 - \delta) \left[(q_f^*)^T (\alpha - \text{Diag}(\beta) (Q_{-f} + q_f^*)) - c_f^T q_f^* \right] + \delta \pi_f^N, \\ & \quad 0 \leq \mathbf{1}^T q_f \leq \text{CAP}_f, \\ & \quad 0 \leq q_f^* \perp \text{Diag}(\beta) Q_{-f} + 2 \text{Diag}(\beta) q_f^* - \alpha + c_f + \gamma_f \geq 0, \\ & \quad 0 \leq \gamma_f \perp \text{CAP}_f - \mathbf{1}^T q_f^* \geq 0. \end{aligned}$$

Table 4. Numerical results on test problem 3 with $\delta = 0.8$

	Variable	Nash solution	Lower bound	Upper Bound	MPEC
Generation	q_{11}	67.2087	73.7309	73.7094	73.7100
	q_{12}	51.9174	32.2374	32.2556	32.2537
	q_{13}	80.8672	0	0	0
	total	172.6016	105.9683	105.9650	105.9637
	q_{21}	24.5257	35.9905	35.9814	35.9731
	q_{22}	12.8997	0	0	0
	q_{23}	12.5745	0	0	0
	total	50	35.9905	35.9814	35.9731
	q_{31}	67.2087	0	0	0
	q_{32}	51.9241	0	0	0
	q_{33}	80.8672	195.8417	195.8780	195.8765
	total	200	195.7413	195.8780	195.8765
	q_{41}	37.2290	39.2440	39.2354	39.2392
	q_{42}	24.5141	0	0	0
	q_{43}	32.8999	0	0	0
total	94.6430	39.2440	39.2354	39.2392	
q_{51}	41.5989	0	0	0	
q_{52}	28.5095	86.0862	86.0725	86.0721	
q_{53}	39.8916	3.6629	3.6653	3.6665	
total	110	89.7491	89.7378	89.7386	
Profits	firm 1	1044.5708	1275.5888	1275.7010	1275.7203
	firm 2	221.3998	434.8644	434.8677	434.7786
	firm 3	1644.5708	1963.2681	1963.2527	1963.2404
	firm 4	217.5821	395.6882	395.7245	395.7748
	firm 5	690.6684	945.9263	945.7592	945.7821
NBO	$\theta(q)$	0	7.1450994E11	7.145111E11	7.145107E11

In order to utilize an NLP solver, we rewrite the two complementarity constraints in the above MPEC as follows:

$$(q_f^*)^T [\text{Diag}(\beta) Q_{-f} + 2 \text{Diag}(\beta) q_f^* - \alpha + c_f + \gamma_f] \leq 0$$

$$(\gamma_f)^T (\text{CAP}_f - \mathbf{1}^T q_f^*) \leq 0.$$

The resulting nonlinear program remains nonconvex. Yet, SNOPT successfully terminated with a solution reported in the last column of Table 4. Notice that for this test problem, the solutions produced by the lower and upper bounding scheme are very close to that produced by SNOPT, thereby once again demonstrating the effectiveness of the bounding schemes. Note also the closeness of the NBO values produced by the two bounding schemes.

7. Conclusion

In this paper, we have introduced an optimization formulation for the collusive game problem and developed upper and lower bounding procedures for its global solution. We have also obtained numerical results that establish the effectiveness of the bounding procedures in obtaining viable solutions, which are supported by publicly available NLP solvers. As the collusive optimization problem is not shown to be convex in general, the bounding procedures, along with their theoretical properties established in Propositions 5 and 8 and their ability to match the results from the NLP solvers, provide confidence on the high fidelity of the numerical solutions obtained by all algorithms. The next step in our research should be application of the proposed model and algorithms to computation of equilibria for actual markets. This will be reported elsewhere.

Acknowledgements. This work is based on research partially supported by the National Science Foundation under an EPNES grant ECS-0224817. The authors are very grateful to two referees who have made many constructive comments that have helped to improve the presentation and results of the paper.

References

1. Bernheim, B.D., Whinston, M.D.: Multimarket contact and collusive behavior. *RAND J. Economics* **21**, 1–26 (1990)
2. Bonnans, J.F., Shapiro, A.: *Perturbation Analysis of Optimization Problems*. Springer-Verlag, New York, 2000
3. Bunn, D.W., Oliveira, F.S.: Evaluating individual market power in electricity markets via agent-based simulation. *Ann. Oper. Res.* **121**, 57–77 (2003)
4. Danskin, J.M.: The theory of min-max with applications. *SIAM J. Appl. Math.* **14**, 641–664 (1966)
5. Dirkse, S.P., Ferris, M.C.: The PATH solver: A non-monotone stabilization scheme for mixed complementarity problems. *Optimization Methods and Software* **5**, 123–156 (1995)
6. Dubey, P.: Inefficiency of Nash equilibria. *Math. Oper. Res.* **11**, 1–8 (1986)
7. Facchinei, F., Pang, J.S.: *Finite-Dimensional Variational Inequalities and Complementarity Problems*. Springer-Verlag, New York, 2003
8. Fourer, R., Gay, D.M., Kernighan, B.W.: *AMPL: A Modeling Language for Mathematical Programming*. Duxbury Press, Brooks/Cole Publishing Company, 2002
9. Friedman, J.W.: *Oligopoly and the Theory of Games*. North-Holland, Amsterdam, 1977
10. Fudenberg, D., Tirole, J.: *Game Theory*. MIT Press, Cambridge, 1991

11. Harrington, J.E., Jr.: Collusion in multiproduct oligopoly games under a finite horizon. *International Economic Review* **28**, 1–14 (1987)
12. Harrington, J.E., Jr.: The determination of price and output quotas in a heterogeneous cartel. *International Economic Review* **32**, 767–792 (1991)
13. Herrero, M.J.: The Nash program: non-convex bargaining problems. *J. Economic Theory* **49**, 266–277 (1989)
14. Horst, R., Thoai, N.V.: DC programming: Overview. *J. Optimization Theory and Applications* **103**, 1–43 (1999)
15. Kaneko, M.: An extension of the Nash bargaining problem and the Nash social welfare function. *Theory and Decision* **12**, 135–148 (1980)
16. Luo, Z.Q., Pang, J.S., Ralph, D.: *Mathematical Programs With Equilibrium Constraints*. Cambridge University Press, Cambridge, England, 1996
17. Nash, J.F.: The bargaining problem. *Econometrica* **28**, 155–162 (1950)
18. Osborne, M.J., Rubinstein, A.: *Bargaining and Markets*. Academic Press, San Diego, 1990
19. Puller, S.L.: Pricing and firm conduct in California's deregulated electricity market. PWP-080, Power Program, University of California Energy Institute, Berkeley, 2001
20. Rockafellar, R.T., Wets, R.J.-B.: *Variational Analysis*. Springer-Verlag, Berlin, 1998
21. Rothkopf, M.H.: Daily repetition: A neglected factor in the analysis of electricity auctions. *The Electricity J.* **12**, 61–70 (1999)
22. Schmalensee, R.: Competitive advantage and collusive optima. *International J. Industrial Organization* **5**, 351–367 (1987)
23. Sweeting, A.: Market outcomes and generator behaviour in the England and Wales wholesale electricity market, 1995-2000. Department of Economics, Massachusetts Institute of Technology, presented at the 7th Annual POWER Research Conference on Electricity Industry Restructuring, University of California Energy Institute, Berkeley, March 22, 2002
24. Tao, P.D., An, L.T.H.: A d.c. optimization algorithm for solving the trust-region subproblem. *SIAM Journal on Optimization* **8**, 476–505 (1998)
25. Tawarmalani, M., Sahinidis, N.V.: Convex extensions and envelopes of lower semi-continuous functions. *Mathematical Programming, Series A* **93**, 247–263 (2002)
26. Tawarmalani, M., Sahinidis, N.V.: *Convexification and Global Optimization in Continuous and Mixed-Integer Nonlinear Programming: Theory, Algorithms, Software, and Applications*. (Kluwer Academic Publishers, Dordrecht, 2002). [Volume 65 in “Nonconvex Optimization And Its Applications” series.]

Copyright of Mathematical Programming is the property of Springer Science & Business Media B.V.. The copyright in an individual article may be maintained by the author in certain cases. Content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.