Center for Research on Economic and Social Theory
and
Department of Economics
Working Paper Series

A Computational Algorithm for Optimal Collusion in Dynamic Cournot Oligopoly

David Roth

December, 1992
Number 93-11

DEPARTMENT OF ECONOMICS
University of Michigan
Ann Arbor, Michigan 48109-1220

A COMPUTATIONAL ALGORITHM FOR OPTIMAL COLLUSION IN DYNAMIC COURNOT OLIGOPOLY

by DAVID ROTH[1]

Department of Economics
University of Michigan
Ann Arbor, MI 48109

December 1992

Abstract *Recent work in oligopoly theory has focused on the optimal degree of tacit collusion sustainable in a dynamic oligopoly. Such analysis for price–setting oligopoly has advanced substantially; however, analysis of Cournot oligopoly has not proceeded apace. This paper analyzes best and worst* strongly symmetric *equilibria for dynamic quantity–setting oligopoly with demand dynamics, such as business cycles, that are independent of previous quantity choices. Two primary results are obtained. First, the paper shows through construction that best and worst strongly symmetric equilibria exist, and that the worst equilibria may be constructed, as in Abreu* (1983, 1986), *so that they prescribe optimal equilibrium play from periods 2 onward. That is, they have a* stick–and–carrot *structure. Second, it provides a simple algorithm for the computation of the extremal equilibria.*

[1] This paper is based upon chapter 2 of my Ph. D. thesis. I would like to thank my advisor David Pearce for his advice and encouragement. This work was supported by an Anderson Fellowship from the Cowles Foundation, Yale University.

I. INTRODUCTION

The maximal degree of tacit collusion attainable in a dynamic oligopoly with perfect monitoring has been the subject of much recent research.[2] For price–setting oligopoly, there has been a remarkably rich literature, with the focus upon extending the results from purely repeated games to games with demand dynamics such as business cycles.[3] However, the literature on Cournot oligopoly has not proceeded past the analysis of purely repeated interaction. Abreu's work (1983, 1986) is the classic, showing that punishments worse than Cournot–Nash reversion are generally available, and that worst *symmetric* punishments may be chosen to have a stick–and–carrot structure. This paper extends the analysis of extremal equilibria of quantity–setting oligopoly from purely repeated games to games with demand dynamics that are independent of the past play of the game.

Two primary results are obtained. First, the paper shows that best and worst *strongly symmetric* equilibria exist and that the worst equilibria may be constructed, as in Abreu (1983, 1986) so that they prescribe optimal equilibrium play in all periods past the first. That is, the worst equilibria have a stick–and–carrot structure in this broader class of games as well. This result is straightforward and rather unsurprising in light of Abreu's earlier work.

On the other hand, the paper also develops a simple technique for the computation of the extremal equilibria. The computation is based upon incentive–constrained dynamic programming techniques. General approaches to such computation have been developed elsewhere (Abreu, Pearce, and Stacchetti (1986, 1990), Fudenberg, Levine, and Maskin (1991)). These general approaches apply to the present problem, yet are computationally quite burdensome. The technique developed in the present paper is, by contrast, extremely

[2] Related work on collusion in oligopolies with imperfect monitoring includes Porter (1983), Green and Porter (1984), and Abreu, Pearce, and Stacchetti (1986).

[3] See Rotemberg and Saloner (1986), Haltiwanger and Harrington (1991), Kandori (1991).

simple, as it only iterates upon artificial best and worst value functions, not upon entire value sets, during computation.

That simplified computation of this sort may apply to best symmetric equilibria in games with apriori known worst punishments has recently been noted elsewhere (Roth (1992)). There I show that, when the worst equilibrium is known, the best equilibrium can be computed through dynamic programming iterating only upon an artificial best value function, keeping the known worst equilibrium as a threatened punishment. For example, Bertrand oligopoly has a known worst equilibrium of $p=MC$ and zero profits forever; thus, the techniques developed in Roth (1992) can be used to compute the optimal sytmmetric equilibria for all of the above–mentioned papers on dynamic price–setting oligopoly. Here, on the other hand, the computation is of both the best and worst equilibria, with the best equilibrium acting as a promised reward for following a suggested action (in both the best *and* worst equilibria, due to the stick–and–carrot structure), and the worst equilibrium acting as threatened punishment. This structure provides us with the simplicity of computation of extremal strongly symmetric equilibria. I defer details until Section 3.

The remainder of the paper is structured as follows. Section 2 introduces the basic model under consideration. Section 3 provides equilibrium characterization and computation. Section 4 summarizes.

II. THE MODEL

The game to be studied is a dynamic, infinite–period, symmetric quantity–setting oligopoly game of perfect monitoring. In each period, there is a state of demand which is known to the oligopolists. Demand follows a Markov transition: Let $z_t \in Z$ be the state at time t, where Z is a subset of some Euclidean space $\mathbb{R}^M$. There are N identical firms. At time t, firm n chooses a quantity level $q_{nt} \in [0,\infty)$ simultaneously with the other firms. Let

$S_n = [0,\infty)$, $n=1,\dots,N$, be the set of production levels possible for the n^{th} firm, and let $S = \times_{n=1}^{N} S_n$ be the set of quantity profiles for all firms.

The inverse demand function at time t is

$$p_t = p(Q_t; z_t),$$

where $Q_t = \Sigma_{n=1}^{N} q_{nt}$. The single–period profit function for firm n, $\Pi_n : S \times Z \to \mathbb{R}$ is as usual:

$$(1) \qquad \Pi_n(q_1,\dots,q_N; z) = p(\sum_{i=1}^{N} q_i \,; z)\, q_n - c\, q_n,$$

where c is the constant marginal cost common to all firms. $\forall\, q,\, z$, let $\pi(q;z) = \Pi_1(q,\dots,q;z)$. Let us make the following assumptions on demand:

(A1) p is continuous in $(Q;z)$,

(A2) $\exists\, M \in \mathbb{R}$ s.t., $\forall\, z,\, q_1,\dots,q_N$, $\Pi(q_1,\dots,q_N;z) \leq M$,

(A3) $\forall\, z$, $p(0;z) > c$, $\lim_{Q\to\infty} p(Q;z) = 0$, and p is monotonic.

(A1), continuity, is necessary for the existence of well–defined, continuous profit functions and optimal equilibria, and for some concepts to be introduced later to be well–defined. (A2) and (A3) ensure that for each z there is some level of production for each firm $\bar{q}(z)$ such that no production above $\bar{q}(z)$ would ever be played in any equilibrium at the state of demand z, for the large short–term losses (A3) could never be recouped by the bounded present discounted value of future profits (A2). We thus will be able to restrict attention to the *compact* move space $[0,\bar{q}(z)]$ when analyzing equilibrium, and to other compact sets when considering incentive compatibility more broadly.

The dynamics of the game are described by a stationary Markov transition density function, which is independent of the actions taken by the players. Let the density $f(z_{t+1}|z_t)$ give the distribution of z_{t+1} given z_t. Firms know the state when they make their production decisions each period; given the state z, single–period profits are given deterministically by (1).

For analysis of equilibrium, let us make the following definitions:

(D1) $\pi^*(q;z) := \max_{q'} \{\Pi_1(q',q,\ldots,q;z)\}$

(D2) $d(q;z) := \pi^*(q;z) - \pi(q;z)$.

Then $d(q;z)$ is the greatest value gained by deviating for one period from the symmetric action profile in which production level q is chosen by all firms at state z; $\pi^*(q;z)$ is the total one–period profits received by deviating from production level q at state z. Each of $d(\ )$ and $\pi^*(\ )$ exists and is continuous, due to (A1), and the fact that $\pi^*(\ )$ is the solution to profit maximization, then, over a continuously varying residual demand curve that drops below the marginal cost curve.

For the remainder of the paper, I restrict attention to strongly symmetric strategy profiles, that is, strategy profiles which both on and off the equilibrium path prescribe similar behavior for all firms. I will use the notation σ for both strategies and strategy profiles. Let $G^{\infty}(\delta)(z)$ denote the full supergame with discount factor δ and first–period state variable z. Throughout, assume that all past production levels of all firms, and all demand histories are common knowledge. That is, monitoring is perfect and perceptions about demand are correct. Assume that players discount future payoffs at the common factor of δ, and discount to the *end* of period 1 (first period profits are not discounted). Firms are interested in maximizing the expected sum of discounted profits.

III. EQUILIBRIUM

In this section, I analyze best and worst strongly symmetric sequential equilibria (hereafter, simply equilibria) of the supergames $G^{\infty}(\delta)(z)$. First, I show that worst and best equilibrium value functions exist. By this I mean that, for each z, the set $V(z)$ of equilibrium expected present discounted future profits for $G^{\infty}(\delta)(z)$, has a worst and a best element. In proving the existence of worst and best equilibria, I construct a *stick–and–carrot* equilibrium which gives the worst equilibrium value, and a *stationary* equilibrium giving the best value (Theorem 1). That is, the worst equilibrium of $G^{\infty}(\delta)(z)$ has a low–profit quantity played in the first period, and the optimal quantity function played thereafter; the best equilibrium has the optimal quantity function played after all histories.

I then develop an incentive–constrained dynamic programming technique which I employ in the computation of the worst and best equilibria. The limits of this procedure exist (Lemma 1), are equilibrium value functions (Lemmas 2–4), and are (weakly) dominated by and dominate, respectively, the worst and best equilibrium value functions. Thus they are themselves the worst and best equilibrium value functions (Theorem 2). The equilibria giving these value functions can then be extracted and described.

Existence of Worst and Best Equilibria

Define the set of equilibrium values of the game $G^{\infty}(\delta)(z)$ as follows:

(D4) $V(z) := \{\ v \mid \exists$ an equilibrium $\sigma(z)$ of $G^{\infty}(\delta)(z)$ s.t. $v = v(\sigma)\ \}$

where $v(\sigma)$ is the expeted present discounted value of future profits under σ. We call a function $v: Z \to \mathbb{R}$ an *equilibrium value function* if, $\forall\ z$, $v(z) \in V(z)$. This is emphatically *not* the same as saying that there exists an equilibrium such that $v(z)$ is the expected present discounted value of future profits accruing to each firm whenever the state of demand is z. Define then the set of equilibrium value *functions* for notational convenience:

(D5) $V := \{ v{:}Z\to\mathbb{R} \mid \forall z,\ v(z)\in V(z) \}$

Define the limits $\underline{v}^*$, $\bar{v}^*{:}Z\to\mathbb{R}$ of the equilibrium value sets point by point as follows:

(D6) For $z \in Z$, $\underline{v}^*(z) := \inf_{v\in V(z)} v$, and $\bar{v}^*(z) := \sup_{v\in V(z)} v$.

Define the sets of best and worst equilibria of $G^\infty(\delta)(z)$ as follows:

(D7) $\bar{\sigma}(z) := \{\sigma(z) \mid \sigma(z)$ is an equilibrium of $G^\infty(\delta)(z)$, and $v(\sigma(z))=\bar{v}^*(z)\}$, and

(D8) $\underline{\sigma}(z) := \{\sigma(z) \mid \sigma(z)$ is an equilibrium of $G^\infty(\delta)(z)$, and $v(\sigma(z))=\underline{v}^*(z)\}$.

Define a stick–and–carrot equilibrium as follows:

(D9) $\sigma(z)$ is a *stick–and–carrot equilibrium of* $G^\infty(\delta)(z)$ if $\sigma(z)\in\underline{\sigma}(z)$, and, $\forall z'$, $\sigma(z,z')\in\bar{\sigma}(z')$,

where $\sigma(z,z')$ is the equilibrium induced by σ on the remaining periods 2,3,..., of the game, when z' is the state of demand in the second period. That is, $\sigma(z)$ is a stick–and–carrot equilibrium of $G^\infty(\delta)(z)$ if it is a worst equilibrium, but prescribes optimal equilibrium play from periods 2 on. We can then show by construction that worst and best equilibria exist, and that the worst equilibrium value may be attained by a stick–and–carrot equilibrium:

Theorem 1. $\underline{v}^*, \bar{v}^* \in V$. *Furthermore, for each z, there exists a stick–and–carrot equilibrium* $\sigma(z)$ *of* $G^\infty(\delta)(z)$ *such that* $\sigma(z)\in\bar{\sigma}(z)$.

Proof. See Appendix. ||

The proof proceeds by constructing, for $z \in Z$, two sequences of equilibria whose values converge to the infimum and supremum of equilibrium values $\underline{v}^*(z)$ and $\bar{v}^*(z)$. Taking appropriate limits, there are worst and best quantity levels $\underline{q}^*(z)$ and $\bar{q}^*(z)$ for which recursive computation and incentive compatibility requirements hold which make $\underline{v}^*(z)$ and $\bar{v}^*(z)$ equilibrium values:

(2) $\underline{v}^*(z) = \pi(\underline{q}^*(z);(z)) + \delta \int \bar{v}^*(z')\, f(z' \mid z)\, dz'$,

(3) $\bar{v}^*(z) = \pi(\bar{q}^*(z);(z)) + \delta \int \bar{v}^*(z')\, f(z' \mid z)\, dz'$,

(4) $d(\underline{q}^*(z);(z)) \leq \delta \int (\bar{v}^*(z') - \underline{v}^*(z'))\, f(z' \mid z)\, dz'$, and

(5) $d(\bar{q}^*(z);(z)) \leq \delta \int (\bar{v}^*(z') - \underline{v}^*(z'))\, f(z' \mid z)\, dz'$. [4]

We see from (2) that the constructed equilibrium inducing value $\underline{v}^*(z)$ is indeed a stick–and–carrot equilibrium. Note that the approach to proof taken here, emphasizing limits of equilibria, complements the set–valued approach to the existence of best and worst equilibria, which is based upon compactness of the equilibrium value set (Abreu, Pearce, and Stacchetti (1986), 1990), Cronshaw and Luenberger (1990), Fudenberg, Levine, and Maskin (1991)).

Computation

Here I present a method of computing the best and worst equilibrium value

[4] All integrals in the paper should be understood to be taken over the state space Z.

functions which does not rely on computing the entire equilibrium value set. The method is, like its set–valued counterparts, an incentive–constrained analog of the iterative contraction mapping technique of dynamic programming. At its heart is the question of which actions can be supported by an arbitrary promised continuation value with the threat of reverting to another, worse punishment value if malfeasance occurs. The iteration proceeds as follows: A pair of value functions $(\underline{v},\overline{v})$ is input. $\underline{v}$ is treated as a punishment threat, $\overline{v}$ as promised reward *and* continuation value. Given these reward and threat functions, the procedure iterates on $\underline{v}$ $(\overline{v})$ by taking the worst (best) of the actions *incentive compatible* with respect to the punishment–reward pair, and adding these one–period payoffs to the expected discounted continuation value $\delta \int \overline{v}(z')\, f(z'|z)\, dz'$. For $\underline{v}$, $\overline{v}: Z \to \mathbb{R}$, $z \in Z$, make the following definition:

(D10) A production level q is *incentive compatible with respect to* $(\underline{v},\overline{v})$ *at* z if

$$\delta \int (\overline{v}(z') - \underline{v}(z'))\, f(z'|z)\, dz' \geq d(q;z).$$

A production level is thus incentive compatible with respect to the pair $(\underline{v},\overline{v})$ if the discounted expected difference between the reward and the punishment exceeds the one–period benefit of deviation. We can then define the set of effort levels incentive compatible with respect to a value function pair at a particular state of demand as follows:

(D11) $IC(\underline{v},\overline{v})(z) := \{\, q \mid q$ is incentive compatible with respect to $(\underline{v},\overline{v})$ at $z \,\}$

Notice that, since cost and demand are continuous (by (A1)), d is well–defined and continuous, and profits become arbitrarily negative with high output but are bounded (A2, A3), we have:

Proposition 1. *The set* $IC(\underline{v},\overline{v})(z)$ *is compact,* $\forall\, z$, $\forall\, \underline{v}$, $\overline{v}: Z \to \mathbb{R}$.

We can now define the transformations $\underline{B}$, $\overline{B}: C \times C \to C$ as follows:

(D12)
$$\underline{B}(\underline{v},\overline{v})(z) := \min_{q \in IC(\underline{v},\overline{v})(z)} \pi(q;z) + \delta \int \overline{v}(z')\, f(z'|z)\, dz' \text{ , and}$$
$$\overline{B}(\underline{v},\overline{v})(z) := \max_{q \in IC(\underline{v},\overline{v})(z)} \pi(q;z) + \delta \int \overline{v}(z')\, f(z'|z)\, dz' \text{ .}$$

Let $B := (\underline{B},\overline{B})$. To reiterate, the map B chooses the worst and best of the values obtainable by producing in the present period, with continuation values $\overline{v}$, subject to the incentive compatibility constraints given by the punishments and rewards $\underline{v}$ and $\overline{v}$. Since $IC(\underline{v},\overline{v})(z)$ is compact, $\forall\, z$, and $\pi(q;z)$ is continuous in q, the map B is well–defined. Notice that $\underline{B}$ has a stick–and–carrot *structure*, giving the lowest incentive compatible profits today and giving high continuation values, while $\overline{B}$ has the structure of an optimal play, giving both the highest supportable profits today and high continuation.

We then take artificially low and artificially high value functions $\underline{v}_0$, $\overline{v}_0$, and iterate through the map B. Start with a $\overline{v}_0$ which is everywhere at least the greatest possible one–period payoffs forever: For example, set $\overline{v}_0 = M/(1-\delta)$ everywhere, where M is one firm's share of the total monopoly profits for the state of demand z which gives highest monopoly profits. Then $\overline{v}_1 < \overline{v}_0$. Also start with a $\underline{v}_0$ which is everywhere lower than the value possible by getting the lowest single period payoffs in $IC(\underline{v}_0,\overline{v}_0)(z)$ plus the expected continuation values $\delta \int \overline{v}_0(z')\, f(z'|z)$. For greatest simplicity, let $\underline{v}_0(z) = 0$. Then $\underline{v}_1 > \underline{v}_0$, since no production level will be incentive compatible with resect to $(0, M/(1-\delta))$ which gives greater than $\delta M/(1-\delta)$ as losses in the first period, since the firm could always shut down and receive total profits of zero. For all $t \geq 0$, let $\underline{v}_{t+1} = \underline{B}(\underline{v}_t,\overline{v}_t)$, and $\overline{v}_{t+1} = \overline{B}(\underline{v}_t,\overline{v}_t)$. Then $\overline{v}_{t+1} \leq \overline{v}_t$, and $\overline{v}_{t+1} - \underline{v}_{t+1} \leq \overline{v}_t - \underline{v}_t$, $\forall\, t$. Since $\{\overline{v}_t\}$ is now a decreasing, bounded sequence of functions, its (pointwise) limit exists; and since $\{\overline{v}_t - \underline{v}_t\}$ is decreasing and thus has a limit, $\{\underline{v}_t\}$ must also have a limit:

Lemma 1. *Set* $\bar{v}_0 = M/(1-\delta)$ *and* $\underline{v}_0 = 0$ *everywhere, and for all* $t\geq 0$, *let* $\underline{v}_{t+1}=\underline{B}(\underline{v}_t,\bar{v}_t)$, *and* $\bar{v}_{t+1}=\bar{B}(\underline{v}_t,\bar{v}_t)$. *Then* $\bar{v}_\infty := \lim_{t\to\infty} \bar{v}_t$ *and* $\underline{v}_\infty := \lim_{t\to\infty} \underline{v}_t$ *exist.*

Proof. See Appendix. ||

The limit value function pair that we have reached is actually a fixed point of the mapping B, as seen by the following result. Define $F=\{(\underline{v},\bar{v})|B(\underline{v},\bar{v})=(\underline{v},\bar{v})\}$ to be the set of fixed points of B. A straightforward limit argument gives us the following:

Lemma 2. $(\underline{v}_\infty,\bar{v}_\infty)\in F$.

Proof. See Appendix. ||

Since $(\underline{v}_\infty,\bar{v}_\infty)\in F$, $\bar{v}_\infty$ is the best value function possible given itself as continuation value and subject to incentive compatibility requirements based on $(\underline{v}_\infty,\bar{v}_\infty)$ as the punishment–reward pair. Somewhat differently, $\underline{v}_\infty$ is the worst value function given $\bar{v}_\infty$ as continuation and supported by the punishment–reward pair $(\underline{v}_\infty,\bar{v}_\infty)$. Therefore, since it is a value function pair induced by actions supportable by the pair itself, it must be a pair of SSE value functions, as must any other fixed pair of B. The following results confirm this intuition:

Lemma 3. *If* $(\underline{v},\bar{v})\in F$, *then* $\underline{v}$, $\bar{v}\in V$.

Proof. See Appendix. ||

The point of Lemma 3 is that a value function pair which is a fixed point of B has obvious

associated SSE action functions, taken directly from the recursive definition $(\underline{v},\bar{v})=B(\underline{v},\bar{v})$ of the value functions. In particular, there must be incentive compatible output levels which provide the first–period profits which, when added to the discounted expected value of the best payoffs thereafter, give the worst and the best total expected profits. Now, since $(\underline{v}_\infty,\bar{v}_\infty)$ is a fixed point of B (by Lemma 2), it is a pair of SSE value functions (by Lemma 3):

Lemma 4. $\underline{v}_\infty$, $\bar{v}_\infty \in V$.

We may now show that these limits are in fact the worst and best equilibrium value functions:

Theorem 2. $\bar{v}_\infty = \bar{v}^*$, *and* $\underline{v}_\infty = \underline{v}^*$.

Proof. See Appendix. ||

The proof proceeds by showing that at every iteration of the procedure, the difference $\bar{v}_t-\underline{v}_t$ is at least as great as the difference between the best and worst SSE values $\bar{v}^*-\underline{v}^*$. Thus the same weak inequality holds in the limit, making $\bar{v}_\infty$ and $\underline{v}_\infty$ the best and worst SSE value functions themselves. Formally, the difference between $\bar{v}_{t+1}(z)$ and $\underline{v}_{t+1}(z)$ is simply $\max_{q\in IC(\underline{v}_t,\bar{v}_t)(z)} \pi(q;z) - \min_{q\in IC(\underline{v}_t,\bar{v}_t)(z)} \pi(q;z)$, which is at least as great as the difference $\bar{v}^*(z)-\underline{v}^*(z)$ as long as $IC(\underline{v}^*,\bar{v}^*)(z)\subseteq IC(\underline{v}_t,\bar{v}_t)(z)$. This is true if $\bar{v}_t-\underline{v}_t\geq\bar{v}^*-\underline{v}^*$. But $\bar{v}_0$, $\underline{v}_0$ were chosen so that $\bar{v}_0-\underline{v}_0\geq\bar{v}^*-\underline{v}^*$. Thus, by induction, $\bar{v}_t-\underline{v}_t\geq\bar{v}^*-\underline{v}^*$, $\forall\, t$. Thus in the limit $\bar{v}_\infty-\underline{v}_\infty\geq\bar{v}^*-\underline{v}^*$. Thus $\bar{v}_\infty$ and $\underline{v}_\infty$ are the best and worst SSE value functions.

Finally, given the limit functions $\bar{v}_\infty$ and $\underline{v}_\infty$, it is easy to extract equilibria of any $G^\infty(\delta)(z)$ which induce values $\bar{v}_\infty(z)$ and $\underline{v}_\infty(z)$. Lemma 3 shows how to do this, since

$(\underline{v}^*, \bar{v}^*) \in F$. Notice that the equilibrium value functions and corresponding equilibria satisfy the value recursion and incentive compatibility conditions (2)–(5).

The proof rests on the fact that, at each step of the iterative process, at least as much incentive power is provided by the artificial continuation values, in terms of choosing incentive compatible output levels, as is provided by the pair of best and worst equilibria themselves. We begin the scheme with a pair $(\underline{v}_0, \bar{v}_0)$ with this property, and the property is preserved by the map B. Note that were we to start with initial artificial best and worst values which provided less incentive power, then the best and worst equilibrium quantities might never become incentive compatible. This can be seen most starkly as follows. Suppose that for all states of demand z, there exists exactly one one–shot Cournot Nash output level. Further, suppose that we began with each of $\underline{v}_0$ and $\bar{v}_0$ equal to the value of playing each of the respective Cournot Nash levels forever. Then there is no incentive power in the punishment–reward pair $(\underline{v}_0, \bar{v}_0)$, only the Cournot Nash levels will be incentive compatible, $(\underline{v}_0, \bar{v}_0)$ will be a fixed point of B, and thus the iterative procedure will never reach the best and worst equilibrium values. Notice that B is in no relevant sense a contraction mapping and that, as the present example of Cournot Nash payoffs shows, may have multiple fixed points. The precise choice of artificially high and low initial values suggested in Lemma 1 and carried throughout the analysis is fundamental, in contrast to the unimportance of the starting point in ordinary dynamic programming problems.

IV. CONCLUSION

This paper has been concerned with the analysis of best and worst equilibria for dynamic quantity–setting oligopoly. In particular, it has focused on the existence, characterization, and computation of best and worst strongly symmetric equilibria and their associated value functions. Best and worst equilibria were proven to exist by construction; the equilibrium constructed which gave the worst value had a stick–and–carrot structure. Finally, the best and worst equilibria were computed as the limits of an incentive–constrained iterative procedure suggested by earlier dynamic–programming approaches to the study of dynamic games.

The substantial advantage enjoyed by the present approach over set–valued iterative approaches is in computation. The class of dynamic Cournot games studied here has best and worst equilibria which have the best equilibria as continuations; thus computation may be restricted to best and worst values, and may disregard any intermediate values. The drawback of the approach is its limited applicability. Games with stick–and–carrot worst punishments and games with convex equilibrium value sets may be the only games whose worst punishments are unknown a priori to afford such simplified computation. General dynamic games have worst punishments, during whose paths continuation values are not necessarily extreme points of the equilibrium value sets. Complete set–valued approaches to the computation of best and worst equilibria are then in order.

Although the class of oligopolies to which this analysis applies is substantial, we should note that extremely important dynamic elements must be left out as well. Capacity expansion, advertising, product development and location, firm–specific learning–by–doing or research and development, and any other firm–specific lasting investment, must be left out due to asymmetries. Industry–wide lasting investment (e.g., lobbying, advertising, R&D, learning–by–doing, stochastic or not), even when it affects all firms equally, may not be analyzed with the present techniques because stick–and–carrot punishments may no longer be optimal when actions affect future payoffs. Finally, as emphasized also by Abreu, there generally exist asymmetric punishment paths which are worse than the worst symmetric punishments for particular players, and thus the best symmetric equilibrium is generally better than the best strongly symmetric equilibrium. The present techniques, however, only apply to the latter.

APPENDIX

Proof of Theorem 1. The proof proceeds by constructing, for each z, two sequences of equilibria whose limit values are the limits $\bar{v}^*(z)$ and $\underline{v}^*(z)$, which are then shown to be equilibrium values themselves, by construction of particular limit equilibria.

For $z \in Z$, let $\{\bar{v}_{t,z}(z)\}$ be a weakly increasing sequence, with $\bar{v}_{t,z}(z) \in V(z)$, $\forall\, t$, and $\lim_{t\to\infty} \bar{v}_{t,z}(z) = \bar{v}^*(z)$. Then for each t, $\exists$ an action $\bar{q}_{t,z}(z) \in S_1$, and equilibrium value functions $\bar{w}_{t,z}$, $\underline{w}_{t,z}: Z \to \mathbb{R}$ with $\bar{w}_{t,z}(z)$, $\underline{w}_{t,z}(z) \in V(z)$, $\forall\, z$, $\forall\, t$, such that:

(1) $\quad \bar{v}_{t,z}(z) = \pi(\bar{q}_{t,z}(z);z) + \delta \int \bar{w}_{t,z}(z')\, f(z'|z)\, dz' \qquad$ (value recursion), and

(2) $\quad d(\bar{q}_{t,z}(z);z) \le \delta \int (\bar{w}_{t,z}(z') - \underline{w}_{t,z}(z'))\, f(z'|z)\, dz' \qquad$ (incentive compatibility).

Consider the limiting behavior of (1). By construction, $\lim_{t\to\infty} \bar{v}_{t,z}(z) = \bar{v}^*(z)$. Next, since the sequence $\{\bar{q}_{t,z}(z)\}$ is bounded, it has a convergent subsequence. Call its limit $\bar{q}^*(z)$. Restrict attention from here on to this subsequence. By continuity of π, then, $\lim_{t\to\infty} \pi(\bar{q}_{t,z}(z);z) = \pi(\bar{q}^*(z);z)$. Finally, $\lim_{t\to\infty} \int \bar{w}_{t,z}(z')\, f(z'|z)\, dz' = \int \bar{v}^*(z')\, f(z'|z)\, dz'$. Suppose not; then it doen't converge from below. Thus $\exists\, t$, $\epsilon_1 > 0$ s.t. $\bar{v}_{t,z}(z) > \bar{v}^*(z) - \epsilon_1$ since $\{\bar{v}_{t,z}(z)\}$ converges to $\bar{v}^*(z)$, but $\int \bar{w}_{t,z}(z')\, f(z'|z)\, dz' < \int \bar{v}^*(z')\, f(z'|z)\, dz' - \epsilon_1/\delta$. Thus $\exists\, \epsilon_2 > 0$ s.t. $\pi(\bar{q}_{t,z}(z);z) > \bar{v}^*(z) - \delta \int \bar{v}^*(z')\, f(z'|z)\, dz' + \epsilon_2$. But then we can find an equilibrium value function $\bar{w}'_{t,z}$ with integral close enough to that of $\bar{v}^*$ that the equilibrium value $\bar{v}'_{t,z}(z) = \pi(\bar{q}_{t,z}(z);z) + \delta \int \bar{w}_{t,z}(z')\, f(z'|z)\, dz' > \bar{v}^*(z)$, a contradiction. Summarizing these limit results, we have:

(3) $\quad \bar{v}^*(z) = \pi(\bar{q}^*(z);z) + \delta \int \bar{v}^*(z')\, f(z'|z)\, dz'$.

An almost identical proof holds for the construction of a recurisve definition for $\underline{v}^*(z)$. The only difference is that we can make the analogous $\int \bar{u}_{t,z}(z')\, f(z'|z)\, dz'$ converge to $\int \bar{v}^*(z')\, f(z'|z)\, dz'$ by a stick–and–carrot argument instead: Any equilibrium value function can be achieved by a worse first–period payoff and better continuation values, if such are available. Formally, suppose that for some t, $\exists\, \epsilon > 0$ s.t. $\int (\bar{v}^*(z') - \bar{u}_{t,z}(z'))\, f(z'|z)\, dz' > \epsilon$. Then $\exists$ an equilibrium value function $\bar{u}'_{t,z}$ s.t. $\int (\bar{u}'_{t,z}(z') - \bar{u}_{t,z}(z'))\, f(z'|z)\, dz' > 0$. Increase $\underline{q}_{t,z}(z)$ to $\underline{q}'_{t,z}(z)$ so that $\underline{v}_{t,z}(z) = \pi(\underline{q}'_{t,z}(z);z) + \delta \int \bar{u}'_{t,z}(z')\, f(z'|z)\, dz'$; we know we can do this by (A2). Do this for each t in a way such that $\lim_{t\to\infty} \int (\bar{v}^*(z') - \bar{u}'_{t,z}(z'))\, f(z'|z)\, dz' = 0$. Now it just remains to show that each $\underline{q}'_{t,z}(z)$ is incentive compatible. First note that $\pi^*(\underline{q}'_{t,z}(z);z) \le \pi^*(\underline{q}_{t,z}(z);z)$, since it is the monopoly profit for a downward–shifted residual demand curve. We now have:

$$
\begin{aligned}
d(\underline{q}'_{t,z}(z);z) &= \pi^*(\underline{q}'_{t,z}(z);z) - \pi(\underline{q}'_{t,z}(z);z) \\
&\le \pi^*(\underline{q}_{t,z}(z);z) - \pi(\underline{q}'_{t,z}(z);z) \\
&= \pi(\underline{q}_{t,z}(z);z) + d(\underline{q}_{t,z}(z);z) - \pi(\underline{q}'_{t,z}(z);z) \\
&\le \pi(\underline{q}_{t,z}(z);z) - \pi(\underline{q}'_{t,z}(z);z) + \delta \int (\bar{u}_{t,z}(z') - \underline{u}_{t,z}(z'))\, f(z'|z)\, dz' \\
&= \delta \int (\bar{u}'_{t,z}(z') - \underline{u}_{t,z}(z'))\, f(z'|z)\, dz'.
\end{aligned}
$$

This completes this section of the proof. Now take a convergent subsequence of $\{\underline{q}'_{t,z}(z)\}$, with limit $\underline{q}^*(z)$. The analogs of (1)–(3) for $\underline{v}^*(z)$ are then the following:

(1′) $\quad \underline{v}_{t,z}(z) = \pi(\underline{q}_{t,z}(z);z) + \delta \int \bar{u}_{t,z}(z')\, f(z'|z)\, dz' \qquad$ (value recursion), and

(2′) $\quad d(\underline{q}_{t,z}(z);z) \le \delta \int (\bar{u}_{t,z}(z') - \underline{u}_{t,z}(z'))\, f(z'|z)\, dz' \qquad$ (incentive compatibility).

$$(3') \quad \underline{v}^*(z) = \pi(\underline{q}^*(z);z) + \delta \int \bar{v}^*(z')\, f(z'|z)\, dz'.$$

Now it only remains to check that there are equilibria which give $\bar{v}^*(z)$ and $\underline{v}^*(z)$ as values. (3) and (3′) imply that $\bar{v}^*(z)$ is induced by the path $(\bar{q}^*(z),\bar{q}^*,\ldots,\bar{q}^*,\ldots)$, while $\underline{v}^*(z)$ is induced by $(\underline{q}^*(z),\bar{q}^*,\ldots,\bar{q}^*,\ldots)$. By (2), $d(\bar{q}_{t,z}(z);z) \le \delta \int (\bar{w}_{t,z}(z') - \underline{w}_{t,z}(z'))\, f(z'|z)\, dz'$; redefine so that its integral decreases as follows: Let $\underline{w}'_{t,z}(z) = \min_{s \le t} \underline{w}_{s,z}(z')$. Then $\underline{w}'_{t,z}(z) \in V(z)$, and (2) becomes $d(\bar{q}_{t,z}(z);z) \le \delta \int (\bar{w}_{t,z}(z') - \underline{w}'_{t,z}(z'))\, f(z'|z)\, dz'$. Taking limits, we have:

$$\begin{aligned} d(\bar{q}^*(z);z) &= \lim_{t\to\infty} d(\bar{q}_{t,z}(z);z) \qquad \text{(by continuity of } d(\,)) \\ &\le \lim_{t\to\infty} \delta \int (\bar{w}_{t,z}(z') - \underline{w}'_{t,z}(z'))\, f(z'|z)\, dz' \\ &\le \delta \int (\bar{v}^*(z') - \underline{v}^*(z'))\, f(z'|z)\, dz'. \end{aligned}$$

Analogously, $d(\underline{q}^*(z);z) \le \delta \int (\bar{v}^*(z') - \underline{v}^*(z'))\, f(z'|z)\, dz'$. Thus $\bar{q}^*(z)$ and $\underline{q}^*(z)$ are supported by the punishment–reward pair $(\underline{v}^*,\bar{v}^*)$, and the following are equilibria with induced value functions $\underline{v}^*$, $\bar{v}^*$: For $\underline{v}^*$, play the path $(\underline{q}^*,\bar{q}^*,\ldots,\bar{q}^*,\ldots)$ until someone deviates, and then start the equilibrium over again; for $\underline{v}^*$, play the path $(\bar{q}^*,\ldots,\bar{q}^*,\ldots)$ until a deviation occurs, and then play the equilibrium giving $\underline{v}^*$. Finally, the worst equilibrium constructed is indeed a stick–and–carrot equilibrium. ||

Proof of Lemma 1. The proof proceeds by induction. We wish to show that if $\bar{v}_{t+1} \le \bar{v}_t$, and $\bar{v}_{t+1}-\underline{v}_{t+1} \le \bar{v}_t-\underline{v}_t$, then $\bar{v}_{t+2} \le \bar{v}_{t+1}$, and $\bar{v}_{t+2}-\underline{v}_{t+2} \le \bar{v}_{t+1}-\underline{v}_{t+1}$. But:

$$\bar{v}_{t+2}(z) = \bar{B}(\underline{v}_{t+1},\bar{v}_{t+1})(z)$$

$$\begin{aligned} &= \max_{q \in IC(\underline{v}_{t+1},\bar{v}_{t+1})(z)} \pi(q;z) + \delta \int \bar{v}_{t+1}(z')\, f(z'|z)\, dz' \\ &\le \max_{q \in IC(\underline{v}_t,\bar{v}_t)(z)} \pi(q;z) + \delta \int \bar{v}_{t+1}(z')\, f(z'|z)\, dz' \end{aligned}$$

since $\bar{v}_{t+1}-\underline{v}_{t+1} \le \bar{v}_t-\underline{v}_t$ implies $IC(\underline{v}_{t+1},\bar{v}_{t+1})(z) \subseteq IC(\underline{v}_t,\bar{v}_t)(z)$, $\forall\, z$,

$$\le \max_{q \in IC(\underline{v}_t,\bar{v}_t)(z)} \pi(q;z) + \delta \int \bar{v}_t(z')\, f(z'|z)\, dz',$$

since $\bar{v}_{t+1} \le \bar{v}_t$, and f is nonnegative,

$$= \bar{B}(\underline{v}_t,\bar{v}_t)(z) = \bar{v}_{t+1}(z),$$

Next,

$$\begin{aligned} \bar{v}_{t+2}(z)-\underline{v}_{t+2}(z) &= \bar{B}(\underline{v}_{t+1},\bar{v}_{t+1})(z)-\underline{B}(\underline{v}_{t+1},\bar{v}_{t+1})(z) \\ &= \max_{q \in IC(\underline{v}_{t+1},\bar{v}_{t+1})(z)} \pi(q;z) - \min_{q \in IC(\underline{v}_{t+1},\bar{v}_{t+1})(z)} \pi(q;z) \\ &\le \max_{q \in IC(\underline{v}_t,\bar{v}_t)(z)} \pi(q;z) - \min_{q \in IC(\underline{v}_t,\bar{v}_t)(z)} \pi(q;z) \end{aligned}$$

since $\bar{v}_{t+1}-\underline{v}_{t+1} \le \bar{v}_t-\underline{v}_t$ implies $IC(\underline{v}_{t+1},\bar{v}_{t+1})(z) \subseteq IC(\underline{v}_t,\bar{v}_t)(z)$, $\forall\, z$,

$$\begin{aligned} &= \bar{B}(\underline{v}_t,\bar{v}_t)(z)-\underline{B}(\underline{v}_t,\bar{v}_t)(z) \\ &= \bar{v}_{t+1}(z)-\underline{v}_{t+1}(z). \end{aligned}$$

But, since $\bar{v}_0$ was chosen artificially high and $\underline{v}_0$ artificially low, $\bar{v}_1 < \bar{v}_0$ and $\underline{v}_1 > \underline{v}_0$. Thus $\bar{v}_1-\underline{v}_1 \le \bar{v}_0-\underline{v}_0$. Thus, $\bar{v}_{t+1} \le \bar{v}_t$, and $\bar{v}_{t+1}-\underline{v}_{t+1} \le \bar{v}_t-\underline{v}_t$, $\forall\, t$. Since $\{\bar{v}_t\}$ is a decreasing, bounded sequence, its limit $\bar{v}_\infty$ exists; similarly does the limit of $\{\bar{v}_t-\underline{v}_t\}$ exist, and thus the limit $\underline{v}_\infty$ of $\{\underline{v}_t\}$ must exist as well. ||

Proof of Lemma 2.

$$\bar{B}(\underline{v}_{\infty},\bar{v}_{\infty})(z) = \max_{q\in IC(\underline{v}_{\infty},\bar{v}_{\infty})(z)} \pi(q;z) + \delta\int \bar{v}_{\infty}(z')\, f(z'|z)\, dz'$$

$$= \lim_{t\to\infty} \max_{q\in IC(\underline{v}_{t},\bar{v}_{t})(z)} \pi(q;z) + \delta \lim_{t\to\infty} \int \bar{v}_{t}(z')\, f(z'|z)\, dz'$$

since each of the sets $IC(\underline{v}_t,\bar{v}_t)(z)$ is compact and nonempty, and since $IC(\underline{v}_{t+1},\bar{v}_{t+1})(z) \subseteq IC(\underline{v}_t,\bar{v}_t)(z)$,

$$= \lim_{t\to\infty} \bar{B}(\underline{v}_t,\bar{v}_t)(z)$$
$$= \bar{v}_{\infty}(z).$$

An analogous computation utilizing $\underline{B}$ shows that $\underline{B}(\underline{v}_{\infty},\bar{v}_{\infty})(z)=\underline{v}_{\infty}(z)$. Thus $B(\underline{v}_{\infty},\bar{v}_{\infty})=(\underline{v}_{\infty},\bar{v}_{\infty})$. ||

Proof of Lemma 3. Suppose $(\underline{v},\bar{v})\in F$. Then, define $\underline{q},\bar{q}: Z\to S_1$ as follows: If $\underline{v}(z) = \pi(q;z) + \delta \int \bar{v}(z')\, f(z'|z)\, dz'$, set $\underline{q}(z) = \operatorname{argmin}_{q'}\{d(q';z) | \pi(q';z)=\pi(q;z)\}$. Then if *any* production level giving one–period profits of $\pi(q;z)$ was supported by the pair $(\underline{v},\bar{v})$, $\underline{q}$ must have been, too. Similarly, if $\bar{v}(z) = \pi(q;z) + \delta \int \bar{v}(z')\, f(z'|z)\, dz'$, set $\bar{q}(z) = \operatorname{argmin}_{q'}\{d(q';z) | \pi(q';z)=\pi(q;z)\}$. Then $\underline{v}(z)$ is generated by playing $\underline{q}(z)$ in the first period and following it with the function $\bar{q}$ thereafter, while $\bar{v}(z)$ is generated by playing $\bar{q}(z)$ in the first period and $\bar{q}$ thereafter. But each of $\underline{q}$ and $\bar{q}$ can be supported exactly by the punishment–reward pair $(\underline{v},\bar{v})$. Thus the following is an SSE inducing value function $\underline{v}(z)$: Play the sequence $(\underline{q}(z),\bar{q},\ldots,\bar{q},\ldots)$ until someone deviates, starting the sequence over thereafter. The following SSE induces $\bar{v}$: Play $\bar{q}(z)$, and then the function $\bar{q}$, until someone deviates, playing the SSE inducing $\underline{v}$ thereafter. ||

Proof of Theorem 2. $\underline{v}^*$ and $\bar{v}^*$ are the best and worst equilibrium value functions. We can therefore see directly from their construction in Theorem 1 that $B(\underline{v}^*,\bar{v}^*) = (\underline{v}^*,\bar{v}^*)$.

The proof proceeds by constructing a comparison of the differences between the best and worst equilibrium value functions $\bar{v}^*-\underline{v}^*$ on the one hand, and the limit equilibria $\bar{v}_{\infty}-\underline{v}_{\infty}$ on the other. Suppose that for some t, $\bar{v}_t-\underline{v}_t \geq \bar{v}^*-\underline{v}^*$. Then, for $z\in Z$:

$$\bar{v}_{t+1}(z)-\underline{v}_{t+1}(z) = \max_{q\in IC(\underline{v}_t,\bar{v}_t)(z)} \pi(q;z) - \max_{q\in IC(\underline{v}_t,\bar{v}_t)(z)} \pi(q;z)$$
$$\geq \max_{q\in IC(\underline{v}^*,\bar{v}^*)(z)} \pi(q;z) - \max_{q\in IC(\underline{v}^*,\bar{v}^*)(z)} \pi(q;z),$$

since $\bar{v}_t-\underline{v}_t \geq \bar{v}^*-\underline{v}^*$ implies that $IC(\underline{v}^*,\bar{v}^*)(z) \subseteq IC(\underline{v}_t,\bar{v}_t)(z)$, $\forall\, z$,

$$= \bar{B}(\underline{v}^*,\bar{v}^*)(z) - \underline{B}(\underline{v}^*,\bar{v}^*)(z)$$
$$= \bar{v}^*(z)-\underline{v}^*(z).$$

But $\bar{v}_0$ and $\underline{v}_0$ were chosen artificially high and artificially low, so $\bar{v}_0-\underline{v}_0 \geq \bar{v}^*-\underline{v}^*$. Thus, by induction, $\bar{v}_t-\underline{v}_t \geq \bar{v}^*-\underline{v}^*$, $\forall\, t$. Thus, in the limit, $\bar{v}_{\infty}-\underline{v}_{\infty} \geq \bar{v}^*-\underline{v}^*$. Thus $\bar{v}_{\infty}$ and $\underline{v}_{\infty}$ must be the best and worst SSE value functions themselves. ||