

Sequential Moves And Tacit Collusion: Reaction-Function Cycles In A Finite Pricing Duopoly

KLAUS WALLNER*

July 3, 1998

ABSTRACT. This paper analyzes a finite horizon, sequential move pricing duopoly, restricting attention to Markov strategies. The solution yields stationary patterns, independent of initial conditions, where the reaction-functions follow cycles of three periods. The market price never settles down, and is at all times strictly above marginal cost. Long-run average industry profits are approximately $5/6$ of the monopoly level. These results demonstrate that neither a long horizon, nor non-Markovian strategies are required for persistent profits in dynamic price-competition games.

1. INTRODUCTION

Since their introduction by Bertrand [1], price-competition models have become a powerhouse in economic research. The static simultaneous move Bertrand model predicts marginal cost pricing in a duopolistic market, a result known as 'Bertrand paradox'. While a simple backward induction argument shows that repetition of the static Bertrand game for a finite number of periods simply repeats the 'Bertrand paradox' in each of the periods, various dynamic versions that generate profitable outcomes have been studied.

This paper analyzes a sequential move, finite horizon pricing duopoly. Two firms compete in the price of a homogeneous product, using Markov strategies. Production is costless, and both time and the price-space are discrete. There is no uncertainty, and future profits are not discounted. The firms take turns setting their prices, which means that each firm's current pricing decision locks it in for the following period as well when the rival moves. Demand goes to the firm with the lowest price, and is split equally if the prices are equal.

*This work is based on the first chapter of my Ph.D. dissertation at Columbia University in 1997. I gladly acknowledge the great debt to my supervisor Prajit K. Dutta, who encouraged me to proceed with this research, and provided invaluable input and sorely needed moral support. I am especially grateful to the associate editor, an anonymous referee, and Joergen Weibull for detailed criticisms on all parts of the paper. The constructive comments of Jay P. Choi, Richard E. Ericson, Guido Friebel, Margaret Madajewicz, and seminar audiences at Columbia, The Stockholm School of Economics, Yale, Cornell, Missouri-Columbia, The Norwegian School of Management, The Federal Reserve Bank of New York, and Florida International University are appreciated.

In this game, only the final two periods are 'irregular' since each firm chooses the optimal myopic behavior in its last move of the game; the reaction-functions in all other periods follow three-period cycles. On the equilibrium path of this game, a firm inherits the smallest price above the monopoly level at which the rival is frozen, and begins a cycle by undercutting and making all sales at the monopoly price. The rival then sets a price sufficiently low to deter further undercutting, in fact inducing the first moving firm to raise its price in the last period of the cycle. In response to this relenting the rival begins a new cycle with the same behavior and reversed roles. Long-run average industry profits are approximately 5/6 of monopoly profits.

The related literature on dynamic price competition examines the robustness of the 'Bertrand paradox' in infinite horizon models, and includes games both with sequential (Maskin/Tirole [3]) and simultaneous moves (supergames).¹ Maskin/Tirole [3] shows the existence of both Edgeworth cycle and kinked demand curve equilibria, and they demonstrate long-run profitability in all outcomes. Their Edgeworth cycles entail repeated undercutting from the monopoly price down to marginal cost. Their kinked demand curve outcomes suffer from a multiplicity of possible kinks; A renegotiation-proofness criterion reduces possible kinks to the monopoly price. They conjecture that the results carry over to a finite horizon. The present paper characterizes the outcome yielded by the Markov-restriction in the finite sequential move model, and shows that there are no kinked demand curve equilibria in this case.

Supergame approaches have generated profitable outcomes using strategies that depend on the payoff-irrelevant history of the game. The non-robustness of long-run profitability with respect to a finite horizon is expressed in Tirole ([4], p. 253):

"Collusion cannot be sustained in the supergame framework, even for a long but finite horizon. Thus an unbounded horizon is crucial to the results. This raises the possibility that the results are not robust to finite-length interaction - a reasonable assumption."

The solution in Markov strategies of the present finite model demonstrates that neither a long horizon, nor non-Markovian strategies are required to sustain high profits in dynamic price competition.

2. THE MODEL

Two firms compete in the price of a homogeneous product. The game is played over discrete time periods $t = 1, 2, \dots, 3T + 2$, where $T > 0$ is a finite even integer. The length of the game is common knowledge and there is no uncertainty. The firms take

¹This model does not rely on existing rationales for profitability in static Bertrand models: capacity constraints (or rising marginal costs in general) and differentiated products.

turns choosing their price from a discrete set. Production is costless, and there is no discounting. All demand goes to the lowest price.

The price set in period t by the currently moving firm is p_t , and that set by its rival in the previous period, to which it remains committed in t , is p_{t-1} . A firm cannot change its price in a period in which the rival moves. The market price in t , m_t , is defined as the lesser of the two firms' prices: $m_t = \min\{p_t, p_{t-1}\}$. Let $\pi(m_t)$ denote industry profits at a market price m_t , and let p^M be the monopoly price. The following assumptions are made.

A1: The profit-function $\pi : \mathbb{R}_+ \rightarrow \mathbb{R}$ is three times continuously differentiable with $\pi'' < 0$ and $\pi(0) = 0$, and π achieves its maximum at $p^M > 0$.

By A1 the monopoly price $p^M = \arg \max_{p \in G} \pi(p)$ is a singleton. The following Lemma summarizes implications of the assumption that are useful for the subsequent analysis.

Lemma 1. *Under A1, $\exists \bar{\varepsilon} > 0$, s.t. $\forall \varepsilon \in (0, \bar{\varepsilon})$: i) $\frac{3}{2}\pi(\varepsilon) < \pi(2\varepsilon)$, ii) $\pi(p - \varepsilon) > \frac{3}{4}\pi(p)$, $\forall p > 2\varepsilon$ with $\pi(p) > 0$, iii) $\pi(p^M + \varepsilon) + \pi(p^M - \varepsilon) > \pi(p^M) + \pi(p^M - 2\varepsilon)$, and iv) $\pi(2\varepsilon) < \frac{1}{8}\pi(p^M)$.*

The proof is in the Appendix. The openness-problem of finding the highest undercutting price in a continuous price-space necessitates the use of a discrete price-grid:

A2: The price grid is $G = \{0, \varepsilon, 2\varepsilon, \dots, p^M, \dots, a\varepsilon\}$, with a positive $\varepsilon < \bar{\varepsilon}$, $\bar{\varepsilon}$ as defined in the Lemma, and $a = \max\{f \in \mathbb{N} : \pi(f\varepsilon) \geq 0\} < \infty$.

The properties of the profit function that are the subject of the Lemma hold for all sufficiently fine price grids. The essence of the fineness requirement on the grid is that marginal undercutting does not lower industry profits so much as to render market sharing at the next higher price more profitable. We adopt an arbitrary splitting rule for the case where both firms set equal prices:

A3: If both firms charge the same price, they share the market equally.

Equal splitting is not crucial, so long as matching results in less than total sales in the market.² The last assumption specifies a tie-breaking rule in the case when a firm is indifferent between several prices, all of which maximize its profit stream over the remainder of the game.

A4: If there are multiple best replies to the rival's continuation strategy, a firm selects $p^M + \varepsilon$ if it is a best reply, and the highest best reply otherwise.

²It is easily verified that if no sharing takes place at equal prices, i.e. if then the last moving firm makes all sales, the undercutting incentive disappears and the monopoly price can be sustained in equilibrium.

It will be shown that such indifference regularly occurs when a firm raises its price, since all relenting prices above p^M are undercut to p^M and therefore yield the same payoff stream.

Denote by $\phi(p, p')$ a firm's profit if its current price is p and its rival's price is p' . Under A3,

$$\phi(p, p') = \begin{cases} \pi(p), & \text{if } p < p'; \\ \frac{1}{2}\pi(p), & \text{if } p = p'; \\ 0, & \text{if } p > p'. \end{cases}$$

The set of histories in period t is $H_t = G^t$, an element of which is $h_t \in H_t$; the set of histories of the whole game is $H = \bigcup_{t=1}^{3T+2} H_t$. A history h_t is the observed past sequence of prices, $h_t = (p_0, p_1, \dots, p_{t-1})$ for $t = 1, 2, \dots, 3T + 2$, where p_0 is the arbitrary preexisting price at which the player moving in period 2 is frozen in the opening period. A pure behavior strategy in period t is a mapping from the history of the game to the price grid, $s_t : H_t \rightarrow G$.

Call the player moving in even (odd) periods 'even' ('odd') firm, denoted by superscripts 'e' ('d'). Note that if t is odd, each firm has the same number of moves remaining, otherwise the firm moving in even periods has one more move. The strategy for the odd firm is a collection $s^d = \{s_t\}_{t=1,3,5,\dots}$, and the strategy for the even firm is $s^e = \{s_t\}_{t=2,4,6,\dots}$. The strategy space for the game is $s = s^d \cup s^e$. Correspondingly we divide the histories of the game into sets containing odd and even periods, $H = H^d \cup H^e$. A behavior strategy of the odd firm is $s^d : H^d \rightarrow G$ and of the even firm $s^e : H^e \rightarrow G$. In even periods, the payoff of the even player is $\phi(p_t, p_{t-1})$ and that of the odd player is $\phi(p_{t-1}, p_t)$; in odd periods, the payoff of the even player is $\phi(p_{t-1}, p_t)$ and that of the odd player is $\phi(p_t, p_{t-1})$.

This finite game contains subgames that are uniquely defined by each history h_t . For all histories h_t , (s^e, s^d) form a Nash equilibrium of the subgame starting at history h^t , i.e. $s^e|_{h_t}$ and $s^d|_{h_t}$ are mutually best responses for all h^t . The purpose of introducing the subgame perfect restriction is to narrow down the possibly large set of Nash equilibria by eliminating non-credible threats.

I am further restricting the firms to using *Markov strategies*. A local strategy s_t is a Markov strategy if for all histories h_t and h'_t , $s_t(h_t) = s_t(h'_t)$ whenever $p_{t-1} = p'_{t-1}$. Under the Markov restriction, firms condition their reply solely on the price set by the rival in the preceding period, which is the only payoff relevant part of the history of the game. Such behavior can be interpreted as conventional dynamic reaction functions, i.e. direct responses to the most recent changes in market conditions. Alternatively, it could describe behavior in a situation where the only information consists of the last move of the rival.

The equilibrium concept used in solving this game is that of a *Markov perfect*

equilibrium. A Markov perfect equilibrium is a subgame perfect equilibrium where both players follow Markov strategies. Since we are assuming that each $s_t : H_t \rightarrow G$ is a Markov behavior strategy, we can summarize the local strategies by defining best responses in period t , $B_t : G \rightarrow G$:

$$B_t(p_{t-1}) = s_t(p_0, p_1, \dots, p_{t-1}).$$

We will describe the equilibrium path of the unique Markov perfect equilibrium satisfying the tie-breaking rule A4. This path is characterized by three-period cycles of the form $(3t-2, 3t-1, 3t)$. Suppose for simplicity that $t \geq 2$ is even. The prevailing price of the odd player, p_{3t-3} , will be $p^M + \varepsilon$. The even firm will then set $p_{3t-2} = p^M$, taking the whole market in period $3t-2$. The odd firm then cuts the price, setting $p_{3t-1} = p_t^*$. This price is set just low enough to deter the even firm from undercutting. In fact, the even firm then raises the price to $p^M + \varepsilon$, and the cycle repeats with the even and odd players reversing roles. This behavior is summarized in the following table:

| | Odd Firm's Price | Even Firm's Price | Market Price |
|--------|---------------------|---------------------|--------------|
| $3t-2$ | $p^M + \varepsilon$ | p^M | p^M |
| $3t-1$ | p_t^* | p^M | p_t^* |
| $3t$ | p_t^* | $p^M + \varepsilon$ | p_t^* |

Table 1: Prices over a typical three-period cycle

This cycling continues until the last two periods, when both firms revert to near-monopoly prices. That is,

| | Odd Firm's Price | Even Firm's Price | Market Price |
|--------|------------------|---------------------|---------------------|
| $3T+1$ | p^M | $p^M + \varepsilon$ | p^M |
| $3T+2$ | p^M | $p^M - \varepsilon$ | $p^M - \varepsilon$ |

Table 2: Prices in the endgame

An analysis of the endgame shows that no pricing other than the myopic best response can be sustained in the two last periods of the game. By backward induction, this competitiveness in the endgame has ramifications for the preceding play of the game. In period $3T$ the firm has a choice between selling twice at lower prices and only once (in $3T+2$) at a high price if it relents now. Such a choice recurs regularly as the game gets longer, and key to understanding the behavior of the game is the fact that in response to low prices the relenting option is more profitable.

The full Markov strategies supporting this equilibrium path as a subgame perfect equilibrium are detailed in the proof of Theorem 2 in the Appendix. But the intuition is straightforward. Continuous undercutting is too costly; it pays to relent if the price falls below a certain level, forgoing current sales for the sake of inheriting a higher market price at one's next move two periods later. It is optimal for each firm to

periodically lower to this level and thereby trigger a raise by the rival. The firms take turns with each three-cycle in setting costly strategic traps for each other and then cashing in for a period; this is an ongoing, repetitive battle not for permanent market dominance but for a transitory, short-term strategic advantage. Neither firm can credibly commit to ending the ongoing strategic warfare and maintaining a stable market price.

This behavior is formalized in Markov strategies of the following form:

$$B_{3T+2}(p) = B_{3T+1}(p) = \begin{cases} p^M, & \text{if } p^M < p; \\ p - \varepsilon, & \text{if } \varepsilon < p \leq p^M; \\ \varepsilon, & \text{if } p = \varepsilon, \end{cases}$$

and for each $t = 0, \dots, T$,

$$\begin{aligned} B_{3t}(p) &= \begin{cases} p^M, & \text{if } p^M < p; \\ p - \varepsilon, & \text{if } p_t^* < p \leq p_{t+1}^L \text{ or } p_{t+1}^L + \varepsilon < p \leq p^M; \\ p, & \text{if } p = p_{t+1}^L + \varepsilon; \\ p^M + \varepsilon, & \text{if } p \leq p_t^*; \end{cases} \\ B_{3t-1}(p) &= \begin{cases} p_t^*, & \text{if } p \leq p_t^U \text{ or } p_t^* < p; \\ p - \varepsilon, & \text{if } p_t^U < p \leq p_t^*; \end{cases} \\ B_{3t-2}(p) &= \begin{cases} p^M, & \text{if } p^M < p; \\ p - \varepsilon, & \text{if } p_t^R < p \leq p_t^U \text{ or } p_t^L < p \leq p^M; \\ p_t^U, & \text{if } p \leq p_t^R \text{ or } p_t^U < p \leq p_t^L, \end{cases} \end{aligned}$$

where each three-period cycle of best response functions $(B_{3t-2}, B_{3t-1}, B_{3t})$ is parameterized by four numbers: $(p_t^*, p_t^U, p_t^R, p_t^L)$. These numbers are defined inductively as follows:

- p_T^* is the largest $p \in G$ less than or equal to the unique $p \in [0, p^M]$ solving

$$\pi(p - \varepsilon) + \pi(p - 3\varepsilon) = \pi(p^M - \varepsilon).$$

Then

$$p_t^* = \begin{cases} p_{t+1}^* + \varepsilon, & \text{if } p_{t+1}^* \leq q; \\ r, & \text{if } p_{t+1}^* > q, \end{cases}$$

where q is the smallest $p \in G$ greater than or equal to the unique $p \in [0, p^M]$ solving $\pi(p) = 2\pi(r - \varepsilon)$ and r is the smallest $p \in G$ strictly greater than the unique $p \in [0, p^M]$ solving $\pi(p) = \frac{1}{2}\pi(p^M)$.

- p_t^U is the largest $p \in G$ less than or equal to the unique $p \in [0, p^M]$ solving

$$\pi(p - \varepsilon) = \frac{1}{2}\pi(p_t^*),$$

- p_t^L is the largest $p \in G$ strictly less than the unique $p \in [0, p^M]$ solving

$$\pi(p - \varepsilon) = 2\pi(p_t^U),$$

- p_t^R is the largest $p \in G$ less than or equal to the unique $p \in [0, p^M]$ solving

$$\pi(p - \varepsilon) = \frac{1}{2}\pi(p_t^U).$$

These strategies describe the equilibrium behavior in this game:

Proposition 2. *Best responses of the form $B_{3T+2}(p)$, $B_{3T+1}(p)$ and for each $t = 0, \dots, T$, $B_{3t}(p)$, $B_{3t-1}(p)$ and $B_{3t-2}(p)$ form the unique Markov-perfect equilibrium of the game satisfying assumptions A1-A4.*

The proof of the Proposition is contained in the Appendix. Graphs of the best responses in the last two periods as well as over a typical cycle are shown below (circles indicate the market price along the equilibrium path).

Picture

Picture

Picture

Picture

Picture

Picture

Picture

Picture

Figure 1: The Endgame

Picture

Picture

Picture

Picture

Picture

Picture

Picture

Picture

Figure 2: Steady-State Three-Period Cycle

The parameters of a three-period cycle only depend on the particular p_t^* defined in that cycle. Let the sequence $\{p_t^*\} = (r, r + \varepsilon, \dots, q)$ thus characterize a sequence of three-cycles, and define the integer $m = \frac{q-r+1}{\varepsilon}$. Call this sequence of three-cycles a 'big cycle'. We have the following result concerning the length of the big cycle.

Corollary 3. *Each big cycle defined by $\{p_t^*\} = (r, r + \varepsilon, \dots, q)$ contains $m - 2$ complete three-period cycles and $3m - 2$ periods altogether. The big cycle repeats in identical form forever as $T \rightarrow \infty$.*

Proof. As detailed in the proof of Proposition 1, when a p_t^* reaches the value q , the sequence $\{p_t^*\} = (r, r + \varepsilon, \dots, q)$ switches from q to r instead of continuing to rise by ε ; at this point a new big cycle begins, with the parameters of the first three-cycle defined by $p_t^* = r$. The sequence $\{p_t^*\}$ contains all gridprices between (and including) q and r , i.e. $\frac{q-r+1}{\varepsilon}$ elements, each of which defines the parameters of one three-period cycle. The first and last cycles in the sequence, occurring at the transition between two big cycles, both consist of only two periods, and hence the number of periods in a sequence is $3m - 2$. By the definition of p_t^* , the sequence $\{p_t^*\} = (r, r + \varepsilon, \dots, q)$ repeats as the game lasts longer. ■³

Writing $\Pi(\varepsilon, T)$ for average industry profits in the game of length $3T + 2$ with grid size ε , we have the following Proposition.

Proposition 4. *Average industry profits converge to $\frac{5}{6}\pi(p^M)$ as the game lasts longer; i.e., for a given $\varepsilon > 0$ satisfying A2, $\Pi(\varepsilon, T) \rightarrow \frac{5}{6}\pi(p^M)$ as $T \rightarrow \infty$.*

Proof. Over a cycle, the equilibrium path of market prices consists of one period at p^M followed by two at p_t^* . The prices p_t^* keep moving up by ε with each cycle: $p_t^* = p_{t+1}^* + \varepsilon$. It follows from the definitions of q and r that the profit levels at the prices of this sequence have as approximate lowest and highest values $\frac{1}{2}\pi(p^M)$ and $\pi(p^M)$, respectively. Hence the average profit over a big cycle in the two periods of a three-cycle that have p_t^* as market price is approximately $\frac{3}{4}\pi(p^M)$, which together with p^M in the third period gives approximately $\frac{5}{6}\pi(p^M)$ as average industry profits for large T . ■

Since $r > 0$, the sequence $\{p_t^*\} = (r, r + \varepsilon, \dots, q)$ implies that the market price is strictly bounded away from zero in all periods. Note also that for any given T , $\Pi(\varepsilon, T) \rightarrow \frac{2}{3}\pi(p^M)$ as $\varepsilon \rightarrow 0$. This holds because $p_t^* = p_T^* + \varepsilon(T - t)$ and therefore $\pi(p_t^*) \approx \frac{1}{2}\pi(p^M)$, $\forall t \leq T$.

Since matching is never an optimal strategy over a three-cycle, there does not exist a kinked demand curve equilibrium. Unlike in the infinite horizon model neither is

³When a big cycle ends and another one begins, the regular pattern of repeating complete three-period cycles is interrupted. As detailed in the proof of Proposition 2, the last and first three-period cycle of a big cycle consists of only two periods each, with the third one being skipped over in the transition process. This irregularity interferes with the notation used in describing equilibrium strategies. The timing subscripts in all big cycles from the second onwards ought to take into account the two missing periods per completed big cycle, which would require using the end period rather than the start of the game as reference point. This notationally more complex subscripting is avoided for ease of exposition. The corollary implies that the smaller ε , the longer the first big cycle, for which the timing subscripts are correct, lasts.

perfect collusion sustainable nor in fact is there any stable market price that can persist over time. In the infinite model there is a large number of kinked demand curve equilibria each based on the reciprocal belief that the rival will also match. In the finite model such outcomes are not Markov perfect, since they unravel from the end. Intuitively, near the end, the gain from continued matching gets smaller (as fewer collusive periods remain) relative to the temptation to undercut and capture the whole market for a period. The finite end date acts as 'counter-commitment', i.e. as guarantee to undercut high prices, destroying the credibility of a price-matching strategy.

3. A LINEAR DEMAND EXAMPLE

This section presents the price-path generated by a linear market demand curve of the form $D : q = 1 - \frac{p}{2}$, for $0 \leq p \leq 2$. The associated industry profits are then $\pi = p - \frac{p^2}{2}$, with a profit of $\pi(1) = \frac{1}{2}$ at the monopoly price $p^M = 1$. The distance of the price grid is 0.01.

The following table gives each firm's price, the market price and profits for the final 17 periods of the game, and for the transition phase between the first and second big cycle. Since on the equilibrium path it is never optimal to match the prevailing price of a rival, the profits are both industry profits and at the same time the payoff of the lower-price firm, with the rival making no sales in that period.

| | Odd Firm's Price | Even Firm's Price | Market Price | Profits |
|----------|------------------|-------------------|--------------|----------|
| 3T+2 | 1 | 0.99 | 0.99 | 0.49995 |
| 3T+1 | 1 | 1.01 | 1 | 0.5 |
| 3T | 0.31 | 1.01 | 0.31 | 0.26195 |
| 3T-1 | 0.31 | 1 | 0.31 | 0.26195 |
| 3T-2 | 1.01 | 1 | 1 | 0.5 |
| 3T-3 | 1.01 | 0.32 | 0.32 | 0.2688 |
| 3T-4 | 1 | 0.32 | 0.32 | 0.2688 |
| 3T-5 | 1 | 1.01 | 1 | 0.5 |
| 3T-6 | 0.33 | 1.01 | 0.33 | 0.27555 |
| 3T-7 | 0.33 | 1 | 0.33 | 0.27555 |
| 3T-8 | 1.01 | 1 | 1 | 0.5 |
| 3T-9 | 1.01 | 0.34 | 0.34 | 0.2822 |
| 3T-10 | 1 | 0.34 | 0.34 | 0.2822 |
| 3T-11 | 1 | 1.01 | 1 | 0.5 |
| 3T-12 | 0.35 | 1.01 | 0.35 | 0.28875 |
| 3T-13 | 0.35 | 1 | 0.35 | 0.28875 |
| 3T-14 | 1.01 | 1 | 1 | 0.5 |
| \vdots | \vdots | \vdots | \vdots | \vdots |
| T-177 | 1.01 | 0.9 | 0.9 | 0.495 |
| T-178 | 1 | 0.9 | 0.9 | 0.495 |
| T-179 | 1 | 1.01 | 1 | 0.5 |
| T-180 | 0.91 | 1.01 | 0.91 | 0.49595 |
| T-181 | 0.91 | 0.3 | 0.3 | 0.255 |
| T-182 | 1 | 0.3 | 0.3 | 0.255 |
| T-183 | 1 | 1.01 | 1 | 0.5 |
| T-184 | 0.31 | 1.01 | 0.31 | 0.26195 |
| T-185 | 0.31 | 1 | 0.31 | 0.26195 |
| T-186 | 1.01 | 1 | 1 | 0.5 |

Table 3: Prices and Profits in a Linear Demand Example

In this example, the first (seen from the end) big cycle ends when p^* in period $T - 180$ reaches the value 0.91. This induces the rival in the preceding period to deviate from the regular pattern of undercutting the prevailing price by ε , and instead to lower to the new trigger level of $p^* = 0.30$.⁴ With this trigger level a new big cycle

⁴The value $p^* = 0.30$ results from the condition for p_t^U to become so high that in the preceding period it is optimal for the rival to lower to that level even from p^M or higher prices. Lowering yields two periods of $\pi(p_t^U)$ compared with one of undercutting at p^M : $\pi(p_t^U) > \frac{1}{2}\pi(p^M)$. The threshold level of p_t^U is then the smallest $p \in G$ strictly greater than the unique $p \in [0, p^M]$ solving $p - \frac{1}{2}p^2 = \frac{1}{2}0.5$. The solution to this equation is given by $p = 0.29289$, and the next higher price on the grid is then 0.3.

begins that goes on until once again p^* reaches the value 0.91. Every time this value is reached a new identical big cycle begins as in period $T - 182$. In this example, the repetitive sequence of cycles contains 60 complete three-period cycles (it takes p^* a three-cycle to move up by 0.1, and over a big cycle p^* moves from 0.30 to 0.91), plus a shortened two period transition phase each in the beginning and end of the sequence; there are hence 184 periods contained in the big cycle.⁵

4. CONCLUSION

Finite horizon oligopoly models have long been shunned because it was thought that the end-date introduced persistent irregularities, thereby rendering the models hard and inelegant to solve without adding valuable insight into the nature of competition. Typically one would solve the infinite horizon version of a model, and conjecture that essentially the same solution applies to the same model with a finite, but long horizon. It is hoped that this paper makes finite horizon models more workable by showing that only very few periods in the end are indeed irregular, and that these quickly disappear from the payoff-relevant horizon as the game moves away from the end. A main feature of this model is the cyclical nature of the reaction-functions, which derives from the fact that the firms only have to consider the subsequent two to four periods in finding their profit-maximizing price. The repetitive triggering and relenting pattern over those steady-state cycles offers a new understanding of the nature of dynamic strategic interaction between two rivals.

The second insight gained from this model is that the minimal one-period commitment introduced by sequential order of moves is sufficient to sustain high profitability. Neither a long horizon, nor non-Markovian strategies are necessary for persistent profitability. Unlike in the infinite horizon version of the model, the price never drops to marginal cost, and there are no kinked demand curve equilibria. Countering intuition, the industry is more profitable in the finite horizon compared to the Edgeworth cycles in the infinite case. While preventing a perfectly collusive outcome, the finite end-date also acts as a commitment device for tougher off-path behavior in the Edgeworth cycles, sustaining a higher equilibrium price-path.

This paper considers only the simplest setting. It is not clear how to model sequential moves in a model where more than two firms compete in price. A further research path is to relax the assumption of a fixed sequential order of moves, and endogenize timing and order of moves in a continuous time framework.

⁵In this example the number of periods in the sequence of three-cycles is even, and therefore it is always the even player who moves in a second consecutive triggering period and initiates the transition to a new big cycle. This is not a general feature of the model and depends on the particular demand curve and grid distance chosen.

Note also that the first big cycle starts at a $p_{3T}^* = 0.31$ and hence contains one complete cycle less than the big cycles to follow.

While the present paper shows the non-robustness of Maskin/Tirole [3] with respect to a finite horizon, their results are only demonstrated for discount factors close to one. The infinite horizon game with impatient firms remains a challenge for future research, and possibly the price dynamics obtained here for the finite horizon offer a clue to its solution.

REFERENCES

- [1] J. Bertrand, *Theorie Mathematique de la Richesse Sociale*, J. des Savants (1883), 499-508
- [2] G. Edgeworth, *La Teoria Pura del Monopolio*. Giornale degli Economisti 40 (1897), 13-31. In English: *The Pure Theorie of Monopoly*, in *Papers Relating to Political Economy*, Volume 1, ed. F. Edgeworth, Macmillan, London, 1925.
- [3] E. Maskin and J. Tirole, *A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles*, Econometrica 56 (1988), 571-599
- [4] J. Tirole, *The Theory of Industrial Organization*, MIT Press, Cambridge, MA, 1988
- [5] K. Wallner, *Tacit Collusion: The Solution to a Finite Pricing Duopoly and Evidence from Japanese Non-Life Insurance*, Ph.D. Dissertation, Columbia University, 1997

5. APPENDIX

Proof of Lemma 1. i) Since by strict concavity of π ,

$$\begin{aligned}
 \lim_{\varepsilon \rightarrow 0} \frac{\pi(2\varepsilon)}{\pi(\varepsilon)} &> \lim_{\varepsilon \rightarrow 0} \frac{\pi(\varepsilon) + \varepsilon \pi'(2\varepsilon)}{\pi(\varepsilon)} \\
 &= \lim_{\varepsilon \rightarrow 0} \frac{\pi(\varepsilon)/\varepsilon + \pi'(2\varepsilon)}{\pi(\varepsilon)/\varepsilon} \\
 &= \frac{\pi'(0) + \pi'(0)}{\pi'(0)} = 2,
 \end{aligned}$$

$\exists \bar{\varepsilon}_1 > 0$ such that the claim holds for all positive $\varepsilon < \bar{\varepsilon}_1$. ii) For $p > p^M$, $\pi(p - \varepsilon) > \pi(p)$ and the claim holds trivially. For $p \in (2\varepsilon, p^M]$, since $\lim_{\varepsilon \rightarrow 0} \pi(p - \varepsilon) = \pi(p)$ by continuity of π , for any given $\frac{1}{4}\pi(p)$, $\exists \bar{\varepsilon}_2 > 0$ s.t. $\forall \varepsilon \in (0, \bar{\varepsilon}_2)$, $\pi(p) - \pi(p - \varepsilon) < \frac{1}{4}\pi(p)$. iii) Performing a second-order Taylor-expansion on both sides of the expression, we get for some $k, l \in (0, \varepsilon)$ and $m \in (0, 2\varepsilon)$:

$$\pi(p^M) + \varepsilon \pi'(p^M) + \frac{\varepsilon^2}{2} \pi''(p^M) + \frac{1}{6} \varepsilon^3 \pi'''(p^M + k)$$

$$\begin{aligned}
& +\pi(p^M) - \varepsilon\pi'(p^M) + \frac{\varepsilon^2}{2}\pi''(p^M) - \frac{1}{6}\varepsilon^3\pi'''(p^M - l) \\
& > \pi(p^M) + \pi(p^M) - 2\varepsilon\pi'(p^M) + 2\varepsilon^2\pi''(p^M) - \frac{4}{3}\varepsilon^3\pi'''(p^M - m).
\end{aligned}$$

Simplifying and dividing by ε^2 , and noting that $\pi'(p^M) = 0$ yields

$$0 > \pi''(p^M) + \sigma(\varepsilon),$$

where $\sigma(\varepsilon)$ is a first-order term in ε with $\lim_{\varepsilon \rightarrow 0} \sigma(\varepsilon) = 0$. Letting $\varepsilon \rightarrow 0$, this becomes

$$0 > \pi''(p^M),$$

which always holds by strict concavity of π . Thus $\exists \bar{\varepsilon}_3 > 0$ s.t. $\forall \varepsilon \in (0, \bar{\varepsilon}_3)$, the claim holds. iv) Since $\lim_{\varepsilon \rightarrow 0} \pi(2\varepsilon) = \pi(0) = 0$ and by A1, $\frac{1}{8}\pi(p^M) > 0$, $\exists \bar{\varepsilon}_4 > 0$ s.t. $\forall \varepsilon \in (0, \bar{\varepsilon}_4)$ the claim holds.

Define $\bar{\varepsilon} = \min\{\bar{\varepsilon}_1, \bar{\varepsilon}_2, \bar{\varepsilon}_3, \bar{\varepsilon}_4\}$. Then $\exists \bar{\varepsilon} > 0$ such that the Lemma holds $\forall \varepsilon \in (0, \bar{\varepsilon})$. ■

Proof of Proposition 2. The structure of the proof is as follows. First the optimal behavior in the last two periods of the game is derived. Then it is shown how directly preceding the endgame the first three-period cycle appears. In particular, in $3T$ it is optimal to relent in response to low prices to benefit from a high market price later on; anticipating this, the rival finds it worthwhile to set a low price in $3T - 1$, giving it three consecutive periods of profits. Key to the proof is to show that this situation repeats periodically with alternating roles, generating a sequence of three-cycles. Finally, the case of $p_{t+1}^* > q$ is considered, which brings to an end one big cycle and starts a new, identical one.

Consider the endgame consisting of periods $3T + 2$ and $3T + 1$. In $3T + 2$, the even firm selects the price that maximizes the current period's profit, given the prevailing price of the rival: $p_{3T+2} = \arg \max_{p \in G} \phi(p, p_{3T+1})$. By property ii) of the Lemma, for all prices above 2ε undercutting gives strictly higher profits than matching, and by i) lowering to ε is the best response to 2ε . Only at ε is matching a best response, since the alternatives of undercutting and raising both yield zero profits. In $3T + 1$, the odd firm foresees that a price above ε will not give it any sales in $3T + 2$. It could make two periods of sales by setting $p_{3T+1} = \varepsilon$, but by A1 this is optimal only in response to ε and 2ε , but not to higher prices. In particular, part i) of the Lemma implies that the best response to $p_{3T} = 3\varepsilon$ is 2ε rather than lowering to ε . Hence $B_{3T+1} = B_{3T+2}$.

In periods $3t$, two cases are to be considered. In $t = T$ the even firm can undercut the odd firm, giving it a profit stream of $\pi(p - \varepsilon) + \pi(p - 3\varepsilon)$. Or it could relent

now, forgoing current sales, in order to benefit from the high market price when it captures the whole market in the last period. By A4, the optimal relenting price is $p^M + \varepsilon$, given that the rival will undercut this and any higher price to p^M . The payoff to relenting is then $\pi(p^M - \varepsilon)$ in the last period. The choice between relenting to sell once at a high price and undercutting to sell twice at lower prices determines the cutoff price p_T^* below which relenting is optimal. At prices above $p^M + \varepsilon$, the best response is $p^M + \varepsilon$ rather than p^M by part iii) of the Lemma.

Next, consider periods $3t$, $t < T$. Facing a following three-cycle $3(t+1), 3(t+1)-1, 3(t+1)-2$, the firm moving in $3t$ needs only consider three periods since the payoff stream beyond that is independent of its current action. This is because its own optimal behavior in $3(t+1)-1$ will force the rival to relent in $3(t+1)$. At low prices, the only way to achieve two consecutive periods of profits is to undercut to at most p_{t+1}^R , with the consequence of having to relent to p_{t+1}^* in the next move two periods later. The alternative is to relent now to $p^M + \varepsilon$ (using A4), giving profits of $\pi(p_{t+1}^*)$ in $3(t+1)-1$. Since it follows from the definitions of p_{t+1}^R and p_{t+1}^* that $2\pi(p_{t+1}^R) < \pi(p_{t+1}^*)$, relenting is more profitable than lowering to p_{t+1}^R in response to any prevailing price of the rival. Optimal behavior in $3t$ is then obtained by comparing relenting to undercutting by ε . The former yields $\pi(p_{t+1}^*)$, while for prices above $p_{3t+1}^L + \varepsilon$ the latter yields $\pi(p - \varepsilon) + \pi(p_{t+1}^*)$ and below it, $\pi(p - \varepsilon)$. This choice determines the cutoff price p_t^* above which marginal undercutting is optimal. Note that p_t^* is defined by the equation $\pi(p - \varepsilon) = \pi(p_{t+1}^*)$, and hence $p_t^* = p_{t+1}^* + \varepsilon$.

At $p = p_{t+1}^L + \varepsilon$ matching is more profitable than undercutting since by giving up half the market today the firm avoids having to relent two periods hence: $\frac{1}{2}\pi(p_{t+1}^L + \varepsilon) + \pi(p_{t+1}^*) > \pi(p_{t+1}^L)$. This follows from the fact that for all t , $\pi(p_{t+1}^*) > \frac{1}{2}\pi(p^M)$, which implies $\pi(p_{t+1}^*) > \frac{1}{2}\pi(p_{t+1}^L)$.

To show that $p_{t+1}^L > p_t^*$, we construct from the definitions of p_{t+1}^L , p_{t+1}^U and p_t^* an indirect link between p_{t+1}^L and p_t^* . First note that $2\pi(p_{t+1}^U - \varepsilon) \leq \pi(p_{t+1}^*)$ and $\pi(p_{t+1}^L - \varepsilon) < 2\pi(p_{t+1}^U)$. From the definition of p_{t+1}^L , the latter implies that $\pi(p_{t+1}^L) \geq 2\pi(p_{t+1}^U)$. Second, $p_t^* = p_{t+1}^* + \varepsilon$, and hence $\pi(p_t^*) > \pi(p_{t+1}^*)$. We have

$$\pi(p_{t+1}^L) \geq 2\pi(p_{t+1}^U) > 2\pi(p_{t+1}^U - \varepsilon) \leq \pi(p_{t+1}^*) < \pi(p_t^*).$$

From concavity of π and the fact that the difference between $\pi(p_{t+1}^U)$ and $\pi(p_{t+1}^U - \varepsilon)$ in the first strict inequality is multiplied by 2, it follows that $p_{t+1}^L > p_t^*$.

In periods $3t-1$, at high prevailing prices the firm can either sell twice by lowering to p_t^* , or once by undercutting. The former yields $2\pi(p_t^*)$, which is larger than $\pi(p - \varepsilon)$ from the latter. In response to prices below p_t^* , undercutting gives two periods of profits $2\pi(p - \varepsilon)$. However, up to p_t^U this is less than relenting to p_t^* with a profit of $\pi(p_t^*)$ in the next period. By property ii) of the Lemma, matching is a less profitable reply to p_t^* than undercutting.

The firm moving in periods $3t - 2$ has to consider only two periods, since the continuation strategy of the rival will force it to relent two periods later independent of what it does now. At high prices, it can sell twice by lowering to p_t^U , or once by marginal undercutting. The comparison of the associated payoffs, $2\pi(p_t^U)$ and $\pi(p - \varepsilon)$, defines the cutoff price p_t^L such that for all higher prices marginal undercutting is the best reply. At prices less than p_t^U , undercutting gives $2\pi(p - \varepsilon)$ which for all prices below p_t^R is less than the payoff to relenting to p_t^U , $\pi(p_t^U)$. Since $\pi(p_t^R) > \frac{1}{8}\pi(p^M)$, property iv) of the Lemma ensures that for suitably fine price grids profitable undercutting is still possible at such low prices.

The end of a big cycle occurs when $p_{t+1}^* > q$. Since $p_t^* = p_{t+1}^* + \varepsilon$, and each p_t^U in turn is a positive function of p_t^* , the sequence $\{p_t^U\}_{t=1,2,\dots}$ is rising as seen from the end. From the definition of r , if $p_{t+1}^U \geq r$ the payoff to lowering to p_{t+1}^U , $2\pi(p_{t+1}^U)$, exceeds that of marginally undercutting $p^M + \varepsilon$, $\pi(p^M)$. In that case, what would be regular periods $3(t+1) - 2$ become periods of the type $3t - 1$, with a new $p_t^* = p_{t+1}^U = r$. This r is the smallest price such that selling twice at r is more profitable than one period at p^M . The definition of q obtains from the smallest p_{t+1}^* that defines a $p_{t+1}^U \geq r$. That $q < p^M$ is confirmed by noting that at a price $p_{t+1}^* = p^M - \varepsilon$, p_{t+1}^U is defined as the largest $p \in G$ less than or equal to the unique $p \in [0, p^M]$ solving $2\pi(p - \varepsilon) = \pi(p^M - \varepsilon)$. Concavity of π then implies that $2\pi(p) > \pi(p^M)$, and hence $p \geq r$ which is the condition for the switch to take place. Hence $q < p^M$ as long as $p_{t+1}^U < r$. ■