

DYNAMIC BERTRAND COMPETITION WITH INTERTEMPORAL DEMAND

PRAJIT DUTTA
DEPARTMENT OF ECONOMICS
COLUMBIA UNIVERSITY

ALEXANDER MATROS
DEPARTMENT OF ECONOMICS
UNIVERSITY OF PITTSBURGH

JÖRGEN W. WEIBULL*
DEPARTMENT OF ECONOMICS
STOCKHOLM SCHOOL OF ECONOMICS

March 31, 2004. This revision February 15, 2005.

ABSTRACT. The standard model of dynamic oligopolistic competition views firms as players in a repeated game, where the demand function is the same in every period. This is not a satisfactory model of the demand side if consumers can make intertemporal substitution between periods. Each period then leaves some residual demand to future periods, and pricing in one period may affect consumers' expectations of future prices. In particular, consumers who observe a deviation from collusive firm behavior may anticipate an ensuing punishment phase with lower prices, and may therefore postpone purchases. In a model that incorporates these two additional elements the interaction between the firms no longer constitutes a repeated game. We here develop a simple model of intertemporal demand in a market setting with overlapping cohorts of consumers, and analyze collusive pricing under Bertrand competition. The more patient and forward-looking consumers are and the higher is the rate of consumer turn-over, the easier it is for firms to collude against them.

*This is a substantially revised version of WP 493, Stockholm School of Economics. We thank two anonymous referees and the editor in charge for helpful comments and suggestions. We are also grateful for comments from Cedric Argenton, Guillermo Caruana, Olivier Compte, Avinash Dixit, Roman Inderst, Sergei Izmalkov, Rafael Repullo, and from seminar participants at Boston University, the CEMFI institute in Madrid, Columbia University, University College London and University of Edinburgh, as well as from participants at the Annual Congress of the European Economic Association 2001, at the First Annual International Conference, "Russia 2015: A Long-Term Strategy", 2001, at the 4th Nordic Workshop in Industrial Organization, 2003, in Copenhagen, and at the ESEM annual congress in Stockholm, 2003. Part of Weibull's work was done at Boston University. Matros and Weibull thank the Hedelius and Wallander Foundation for financial support.

1. INTRODUCTION

The Coase conjecture (Coase, 1972) stipulates that a monopolist selling a new durable good cannot credibly commit to the monopoly price, because once consumers have made their purchases at this price, the monopolist will have an incentive to reduce the price in order to capture residual demand from consumers who value the good below the monopoly price. This in turn, Coase claims, would be foreseen by consumers with valuations above the monopoly price, and therefore some of these – depending on their time preference – will postpone their purchase in anticipation of a price fall. Coase’s argument is relevant not only for a monopoly firm in a transient market for a new durable good, but also for oligopolistic firms in a perpetually ongoing market for durable and non-durable goods. If such firms maintain a collusive price above the competitive price under the threat of a price war, as the literature on repeated-games suggests they may, then consumers might foresee price wars in the wake of a defection, and hence not buy from a firm that slightly undercuts the others, but instead postpone purchase to the anticipated subsequent price war. Such dynamic aspects of the demand side runs against the spirit of the usual model of dynamic competition viewed as a repeated-game.¹ Indeed, the interaction is no longer a repeated-game, since the market demand faced by the firms today in general depends on history, both through consumers’ expectation formation and through their residual demand from earlier periods. Consequently, a model with consumers who can make intertemporal substitution between periods falls outside the domain of the standard model of dynamic oligopolistic competition.

In this paper we develop just such a model, one that adds intertemporal economic agents on the demand side to standard Bertrand competition on the firms’ side. We show that, in comparison with the case of a monopoly for a new durable product, the application of Coase’s argument to oligopoly leads to a radically different conclusion: under a wide range of circumstances such intertemporal substitution and foresight on behalf of the consumers facilitates, rather than undermines, monopoly pricing in a recurrent market setting. Our conclusion is, however, in line with the findings in Ausubel and Deneckere (1987) and Gul (1987).² Indeed, there is a literature on the Coase conjecture, building on models of consumers who have the possibility of intertemporal substitution and are endowed with foresight, see for example Gul, Sonnenschein and Wilson (1986) and the just mentioned papers. We here model consumers very much in the same vein. However, while in those models all consumers enter the market in the initial time period, in our model new consumers enter

¹See e.g. Tirole (1988) for repeated-games models of dynamic oligopoly, and Fudenberg and Tirole (1991) for various versions of the folk theorem.

²However, we also show that in some circumstances the effect may go in the same direction as in the Coase conjecture: Collusion may be more difficult if consumers have foresight.

the market in each market period. In this respect we follow Conlisk, Gerstner, and Sobel (1984) and Sobel (1984, 1991). Like here, these authors assume that consumers differ in their valuation of the good and want to buy the good at most once. However, while in those models consumers who have not made a purchase remain forever in the market, and hence residual demand builds up indefinitely over time, our consumers spend a finite (random) time in the market – a fixed fraction of the consumer population leaves the market each period even if they have not made any purchase. Hence, unlike these earlier models, our model allows for the possibility of stationary supply and demand conditions. Moreover, firms in Conlisk, Gerstner, and Sobel (1984) and Sobel (1984, 1991) cannot resist dropping a collusive price, that is, have sales, because the residual demand from consumers with low valuations grows beyond any upper bound. By contrast, firms in our model can sustain the same collusive price in equilibrium. Indeed, most of our analysis is focused on such equilibria. However, we also show that under certain circumstances equilibrium sales are possible also in our model, that is, with a constant consumer population.

More precisely, the market is open over an infinite sequence of market periods. In each period every firm commits to a price in that period. There is a continuum of consumers, and we model this population as consisting of overlapping cohorts, where a new cohort of consumers “are born” (enter) in each period and an equally large set of “old” consumers – that is, who were present in the preceding period – “die” (leave). The birth and death of consumers are assumed to be driven by exogenous factors. The size of the consumer population is thus constant. The population share of newborn consumers in any period is some fixed fraction. This is also the population share of the previous period’s consumer population who died. All consumers have the same probability of dying each period. The life time of every consumer is thus a geometrically distributed random variable with constant hazard rate and finite expectation. The demographic composition of the consumer population is, by contrast, deterministic and stationary. Our model contains the standard repeated-game model as the special case when the whole population is turned over every period.

The good in question is sold in indivisible units, and each consumer wants to buy one unit of the good at most once during his or her lifetime. While having identical life-table distributions, consumers differ as to their valuation of the good. In each period, the newly arrived consumers’ valuations are distributed according to some fixed cumulative distribution function, while the remaining old consumers are divided into two groups: those who already bought a unit, and those who did not yet do so. The valuation distribution in the latter group depends on the history of prices and price expectations.

Following the above-mentioned analyses, we treat firms as players in the game-

theoretic sense but model consumers as price-taking and expectation-forming economic agents with no strategic power. Their aggregate demand constitutes a state variable in a stochastic game played by the firms. The bulk of our analysis is focused on the case of consumers with perfect foresight. However, this paper is not a plea that analysts should always assume all economic agents to have perfect foresight. We believe that consumers and firms may more realistically be modelled as having less than perfect foresight, and we indeed show how our model applies to consumers with behaviorally more plausible expectations. Our position is rather that the contrast in repeated-games models of dynamic oligopolistic competition between, on the one hand, the intertemporal substitution possibilities, sophistication and expectations coordination ascribed to firms, and, on the other hand, the complete lack of intertemporal substitution ascribed to consumers, should be replaced by a milder contrast. Even taking a small step in this direction requires the analyst to go outside the familiar class of repeated-games to the wider and less familiar class of stochastic games. We here outline how such a generalization can be made, and provide some of its most direct implications. We also believe that our model of demand can be a useful work-horse for other dynamic market analyses.

The paper delivers one key trade-off and a few main results, all of which are absent in the standard infinitely repeated Bertrand model. The trade-off has to do with forward-looking consumers' reaction to a (deviant) firm's price-cut. When consumers are forward-looking, a price-cut does not necessarily induce consumers to buy since they might choose to wait for an even lower price in the future; this reduces the attractiveness to firms of deviating from a collusive price. On the other hand, since consumers are long-lived, there are consumers from previous periods who had chosen not to buy at the going high price. Consequently, a price-cut can reel in such consumers; this increases the attractiveness of deviating. The relative size of these two countervailing forces determines whether the payoff to deviation is smaller or larger than in the standard repeated-games model. When consumers are sufficiently patient, the stronger force turns out to be the collusive one. Consumers defer purchases after the first price-cut and wait for the price war to ensue, thus lowering the deviation payoff to undercutting below that in the repeated-games model. By contrast, when consumers are impatient, the fact that there is residual demand from old customers, who add their demand to that of the new high-value ones, implies that deviation payoffs are higher — and collusion is thus harder to sustain — than in the standard repeated-game model.

There is one other interesting consequence of having long-lived consumers. In the standard model, if a firm were to undercut a collusive price, then it would do so by undercutting ever so marginally. The reason for that is that any price below the going price would guarantee the whole market for the deviant, and the industry

revenue is increasing in prices below the going price (assuming that the collusive price does not exceed the monopoly price and that the industry revenue function is single-peaked). In the current model, by contrast, the most profitable under-cutting price may be substantially lower than the collusive price. Such significant under-cutting can be profitable because it attracts old consumers with low valuations. Only a significant price cut can bring in old consumers, since those amongst them who have not yet bought have valuations below the going price.³ Since they represent a positive fraction of the potential buyers, a deviant firm, when the collusive price is at or near the monopoly price, will find it optimal to capture their demand.

Moreover, since the market price in the first period after a unilateral price cut, if anticipated by consumers, will affect their demand in the defection period, the “punishment” of a unilateral price cut not only affects the defector’s future profits but also its profit in the defection period itself. Because of this effect, absent in repeated games, even harsher punishments than grim trigger strategies are possible if the marginal cost is positive, namely, to force the defector to price below marginal cost in the post-deviation period, thus bringing down the profit in the defection period below what it would have been under grim trigger strategies. We identify and analyze a class of such “generalized trigger” strategies, and focus on maximally “harsh” punishments of this sort. Any constant collusive price that can be supported in subgame perfect equilibrium can also be supported by subgame perfect equilibria in such strategies, so this approach allows us to explore the full range of stationary subgame perfect equilibrium outcomes.

The above discussion suggests that it may be profitable for the firms to now and then run a coordinated sale, in equilibrium, and thereby increase their profits above monopoly profits.⁴ By assumption, such sales are anticipated by consumers with perfect foresight and will hence reduce profits in the periods just preceding the sale, but this may be compensated by the profits made during the sale, because of the accumulated residual demand among old low value consumers. During the sale, consumers correctly anticipate reversion to “normal” pricing next period, and hence have no reason to postpone their purchases. This contrasts sharply with the unanticipated price deviations mentioned above, where the demand facing the under-cutting firm is dampened by consumers’ anticipation of an ensuing price war. Such equilibrium sales can be viewed as a form of temporary price discrimination. We provide conditions under which equilibrium sales are profitable/unprofitable.

³Recall that the only old consumers still in the market are those with valuations below or at the going price. At a slightly lower price these consumers have virtually no surplus from buying today but a positive surplus from buying next period (at a low price). So they wait.

⁴Note that the discussion in the previous paragraph referred to an off-equilibrium path phenomenon.

There are of course other models of dynamic Bertrand competition that depart from the repeated-games paradigm. Kirman and Sobel (1974) consider the role of inventories, and Maskin and Tirole (1988) and Wallner (1999) consider the role of alternating moves. Selten (1965a,b) and Radner (1999) introduce consumers who switch suppliers according to observed prices, though not immediately or fully. All these strands of the literature model a dynamic oligopolistic market as a dynamic or stochastic game, though differently from what we do here.

The paper is organized as follows: the model is developed in section 2, and the first steps of the analysis are made in section 3. Our main results concern stationary subgame perfect equilibria are given in section 4. Section 5 analyses equilibrium sales, section 6 briefly discusses consumers with adaptive expectations, and section 7 concludes.

2. THE MODEL

Suppose there are n firms in a market for a homogenous indivisible good. The market operates over an infinite sequence of periods, $t = 0, 1, 2, \dots$. All firms simultaneously set their prices at the beginning of every period and are committed to that price during the period. Let $p_{it} \geq 0$ be firm i 's price in period t . All consumers observe all posted prices, and buy from the firms with the lowest price. The lowest price in any period will be called the *market price* in that period,

$$p_t = \min\{p_{1t}, \dots, p_{nt}\}. \quad (1)$$

If more than one firm asks the market price, then sales are split equally between these. The firms face no capacity constraint and produce the good at a constant marginal cost $c \geq 0$. Hence, each firm's profit in a market period is simply its sales multiplied by the difference between its price and marginal cost. All firms are risk neutral and discount future profits by the same *discount factor* $\delta \in (0, 1)$ between successive market periods. Resale is not possible.

There is a continuum population of consumers, divided into overlapping cohorts. A cohort of new consumers arrive (are "born") each period, and an equally large set of old consumers, that is, consumers who were in the market in the preceding period, exit ("die"). The size of the consumer population is thus constant, and we normalize it to 1. The population share of new consumers in any period is $\alpha \in (0, 1]$, and this is also the share of the previous period's population that exits/dies each period. All consumers have the same probability α of exiting/dying each period. We will call α the consumer *turn-over rate*. It follows that the "life span" S of every individual is a geometrically distributed random variable, with probability α for $S = 1$, probability $(1 - \alpha)\alpha$ for $S = 2$, probability $(1 - \alpha)^2\alpha$ for $S = 3$ etc. The demographic composition of the consumer population, however, is deterministic and

stationary: in all periods the share of newly arrived consumers is α , the share of one period olds is $(1 - \alpha)\alpha$, and, more generally, of s periods olds is $(1 - \alpha)^s\alpha$, for $s = 0, 1, 2, \dots$. We will call the newly arrived *young* and all others *old*.

Each consumer wishes to buy at most one unit of the good during his or her lifetime. Consumers differ as to their valuation v of the good. In each period, young consumers' valuations v are distributed according to a fixed cumulative distribution function $F : \mathbb{R}_+ \rightarrow [0, 1]$. We assume F to be continuous and to be strictly increasing wherever $F(p) < 1$. Let $D : \mathbb{R}_+ \rightarrow [0, 1]$ be defined by $D(p) = 1 - F(p)$. The function D corresponds to the demand function in a static setting. Let the function $\Pi : \mathbb{R}_+ \rightarrow \mathbb{R}$ be defined by $\Pi(p) = (p - c)D(p)$. We will refer to this as the *industry profit function* and assume it to be single-peaked with maximum at some positive price, the *monopoly price*, $p^m = \arg \max_{p \geq 0} \Pi(p)$.

All consumers have the same pure time preference, represented by the discount factor $\gamma \in [0, 1]$: the discounted expected utility from purchase of a unit $\tau \geq 0$ periods later at price $p \geq 0$ is $(v - p)\gamma^\tau$, while the utility of never acquiring the good is normalized to zero.⁵ In view of the probability α of exiting the market or dying, the effective discount factor for purchasing decisions, is $\beta = (1 - \alpha)\gamma$.⁶

2.1. Consumer expectations and choices. Consumers have perfect foresight concerning future prices. In any given period, let p denote the current market price, and let $p^e(\tau)$ be the expected market price τ periods ahead, for $\tau = 0, 1, 2, \dots$ (hence $p^e(0) = p$). For a consumer with valuation v , who has not yet bought a unit, it is optimal to buy in the present period if and only if her utility from doing so is neither exceeded by the utility from never buying the good nor by the expected utility from postponing purchase to some future period, that is, iff

$$v - p \geq \max \{0, \beta[v - p^e(1)], \beta^2[v - p^e(2)], \dots\}. \quad (2)$$

Of particular relevance for the subsequent analysis are scenarios when the price is expected to remain constant in the near future, then drop to a lower price and thereafter not fall any lower. Formally: let σ be a positive integer and suppose $p^e(\tau) = p$ for $\tau = 0, 1, 2, \dots, \sigma - 1$, $p^e(\sigma) = p^e < p$ and $p^e(\tau) \geq p^e$ for all $\tau > \sigma$. The special case $\sigma = 1$ is thus the scenario when “tomorrow’s” price is expected to be

⁵An interesting extension is to allow for consumer heterogeneity also with respect to time preferences.

⁶This parametrization in effect assumes that consumers literally “die” when exiting the population. An alternative scenario, calling for slightly different parametrization, is when exiting consumers migrate to another economy, with other purchasing opportunities — hence, where unspent money has a positive value.

lower than “today’s,” and $\sigma = 2$ the scenario when the price is expected to fall the “day after tomorrow” etc.

Consider such a price scenario, and a consumer with valuation v who has not yet bought the good. Her utility from buying in a “pre-sale” period $\tau < \sigma$ is $\beta^\tau (v - p)$, from buying in the “sales period” σ is $\beta^\sigma (v - p^e)$, and from buying in a “post-sale” period $\tau > \sigma$ is less than or equal to $\beta^{\sigma+1} (v - p^e)$. Hence, it is never optimal to plan to buy in a post-sale period, nor in a pre-sale period other than the current period ($\tau = 0$). The remaining choices are: buy in the current period, in the sales period, or not at all. It is easily verified that the first choice is optimal if and only if the consumer’s valuation v is at least

$$v^\sigma(p, p^e) = \frac{p - \beta^\sigma p^e}{1 - \beta^\sigma}. \quad (3)$$

We will call $v^\sigma(p, p^e)$ the *cut-off valuation level*.⁷ Among the young and those old who have not yet bought a unit, all consumers with valuations $v > v^\sigma(p, p^e)$ will thus buy in the current period, while those with valuations $v < v^\sigma(p, p^e)$ will not buy — they will either wait for the expected price cut or abstain from buying.⁸

In any market period, and under any price expectation scenario $p, p^e(1), p^e(2), \dots$, consider a cohort of consumers who entered s periods ago. From equation (2) it is clear that it is always the upper tail of the cohort’s value distribution that buys. Hence, s periods back, when the cohort was “young,” everybody above some cut-off valuation purchased the good. In the next period, that is, $s - 1$ periods back, some more consumers from the same cohort may have purchased the good at a lower price, and some consumers have “died.” The additional buyers belong to the upper tail among the non-buyers from s periods back, and so on. By the time this cohort reaches the current period, it has shrunk in size by the factor $1 - \alpha$ each period, and there is a highest valuation v_s among the remaining individuals such that all consumers in the cohort with valuation above v_s have already purchased the good, while none of those with lower valuations have done so.⁹ In sum: the size of the cohort that entered s periods ago is $\alpha(1 - \alpha)^s$, and the valuation distribution in that cohort is given by the c.d.f. $F_s(x) = F(x)/F(v_s)$ for $x \in [0, v_s]$.

⁷To see that this decision rule is optimal, note that the utility from current purchase is $v - p$, from purchase in the sales period $\beta^\sigma (v - p^e)$, and from no purchase it is 0. Thus, if $v > v^\sigma(p, p^e)$, the first utility exceeds the second, and, since $p^e \leq p$, it also exceeds the third (zero). If instead $v < v^\sigma(p, p^e)$, then the second utility exceeds the first.

⁸For the sake of definiteness, but without affecting the results (since the valuation distribution is assumed continuous), we assume that consumers who are indifferent between buying now and in the future or not at all, will buy now.

⁹Note also that $v_s \leq v_{s-1}$ for all s ; an older cohort cannot have a higher current cut-off valuation than a younger cohort.

It follows that the value distribution among the old consumers in any period t who have not yet made a purchase is completely described by the vector $v^t = (v_1, v_2, \dots, v_t)$ of all old cohorts' highest valuations. This vector determines, in its turn, current aggregate demand from old consumers, as a function of the current market price and current consumer expectations about future prices. Current demand from young consumers, who make up the population share α , is derived from the original value distribution F .

2.2. Equilibrium. We assume that firms know all past prices announced in all earlier periods, as well as the current aggregate demand function.¹⁰ This information defines the *state* in the stochastic game played by the n firms. A (pure behavior) *strategy* for a firm is accordingly a function that specifies a price to set in each period τ , conditional upon the state in that period. Firms' strategies constitute a *subgame perfect equilibrium* if in all periods and states each firm maximizes its expected discounted future stream of profits, given all other firms' strategies. We note that, since the stage game in a period in general depends on the current state, which in general depends on the price history, the strategic interaction between the firms is not a repeated but a stochastic game.¹¹

3. PRELIMINARIES

Before analyzing the model in full generality, we here pin-point aggregate demand in steady state and examine a special case of the present model that coincides with the usual repeated-games model of dynamic Bertrand competition.

3.1. Aggregate demand. Suppose that in all past periods the market price was p^* and that this price was expected to remain in all future periods.¹² What would current aggregate demand then be if consumers (a) experienced a sudden price cut, $p < p^*$, (b) expected some price $p^e \leq p$ in the next period and (c) expected no future price below p^e ?

Current demand from new consumers would then simply be their population share, α , times the share of new consumers with valuations exceeding their current cut-off

¹⁰In fact, it is sufficient that firms hold correct expectations along the induced price path and after unilateral deviations from this path.

¹¹For a discussion of stochastic - sometimes called Markovian - games, see Fudenberg and Tirole (1991), chapter 12, and see Dutta (1995) for equilibrium characterizations in such games.

¹²More exactly, we here focus on collusion in a market environment where initial conditions have played out their role. We effectively assume an infinite past, or, equivalently, an initial state that is consistent with an infinite past.

valuation, defined in equation (3):

$$D^y(p, p^e) = \alpha D\left(\frac{p - \beta p^e}{1 - \beta}\right). \quad (4)$$

By contrast, current demand from the old would also depend on their past cut-off valuations. When all firms' prices in the past were p^* and were expected to remain at that level, all old consumers with valuations above p^* have already bought a unit, while those with lower valuations have abstained from buying. Hence, we would have $v_s = p^*$ for all cohorts $s \geq 1$. Current demand from the old, whose population share is $1 - \alpha$, would thus be

$$D^o(p^*, p, p^e) = (1 - \alpha) \max\left\{0, D\left(\frac{p - \beta p^e}{1 - \beta}\right) - D(p^*)\right\}. \quad (5)$$

In sum, current *aggregate demand* would be

$$A(p^*, p, p^e) = \alpha D\left(\frac{p - \beta p^e}{1 - \beta}\right) + (1 - \alpha) \max\left\{0, D\left(\frac{p - \beta p^e}{1 - \beta}\right) - D(p^*)\right\}. \quad (6)$$

This shows that, for sufficiently low under-cutting prices p , aggregate demand emanates from both the old and young, while for higher under-cutting prices p it emanates only from the young — old consumers with high valuations have already bought a unit. The intermediate under-cutting price that separates these two cases is a convex combination of current and past expected prices: $\bar{p} = \beta p^e + (1 - \beta)p^*$. The more patient consumers are, the more weight is given to the currently expected price for the next period.

3.2. The repeated-games model. The standard model of dynamic oligopoly corresponds to the special case when a new batch of consumers enter each period, that is, when $\alpha = 1$ and hence $\beta = 0$. Then equation (6) gives

$$A(p^*, p, p^e) = D(p), \quad (7)$$

for all p^* , p and p^e . The oligopoly thus faces the same demand function each period. We will refer to this special case as *the standard repeated-games model*.

Trigger strategies supporting a constant collusive price p^* exceeding marginal cost can then be defined in the usual way: all firms ask the price p^* in the initial period and continue to do so in all future periods as long as no firm undercuts this price. Otherwise, all firms set the price c , their marginal production cost, from that

period on.¹³ Suppose that consumers observe all past and current prices and have perfect foresight concerning future prices (at least on the price path induced by the firms' strategy profile and after any unilateral deviation from this path). Given any current market price p and the price history where all firms asked the price p^* in all past periods, we then have

$$p^e(\tau) = \begin{cases} p^*, & \text{if } p = p^*, \\ c, & \text{if } p \neq p^*, \end{cases} \quad (8)$$

for all $\tau > 0$. Such expectations fall into the class of price scenarios discussed in section 2. Hence, cut-off valuations satisfy equation (3) with $\sigma = 1$.

A trigger-strategy profile, in which all firms quote the same collusive price p^* in all periods until a price deviation is detected, from which time on they all quote the price c , constitutes a subgame perfect equilibrium in this special case if and only if

$$\Pi(p) \leq \frac{\Pi(p^*)}{n(1-\delta)} \quad (9)$$

for all $p < p^*$, where Π is the *industry profit function*, defined in section 2. The quantity on the left-hand side of (9) is the present value of the profit to a firm that undercuts the collusive price by posting a price $p < p^*$ — such a firm will earn zero profit in all future periods — and the quantity on the right-hand side is the present value of the firm's profit were it to remain at the collusive price p^* . By continuity of the value distribution, inequality (9) holds if and only if

$$\max_{p \in [0, p^*]} \Pi(p) \leq \frac{\Pi(p^*)}{n(1-\delta)}. \quad (10)$$

Since the industry profit function Π by hypothesis is single-peaked, the left-hand side in equation (10) is simply $\Pi(p^*)$ if the collusive price p^* does not exceed the monopoly price — a deviating firm then wants to undercut the going price only slightly. Hence, a collusive price $p^* \in (c, p^m]$ is supported by a subgame perfect equilibrium in grim trigger strategies if and only if

$$\delta \geq 1 - 1/n. \quad (11)$$

4. CONSTANT-PRICE COLLUSION

Having considered constant collusive pricing in the special case when $\alpha = 1$ and $\beta = 0$, we now turn to situations when $0 < \alpha < 1$ and $0 < \beta < 1$. The interaction is no

¹³In order to make this a complete and subgame perfect strategy specification, assume all firms set the price c also in case current aggregate demand is inconsistent with p^* having been the going price and price expectation in all preceding periods.

longer a repeated game, and an additional consideration comes into play: the market price in the first period after a deviation from a collusive price, if anticipated by consumers, will affect their demand in the defection period. Hence, the “punishment” that follows upon a unilateral price cut not only affects the continuation payoffs to the defector but also the defector’s payoff in the defection period itself. Because of this effect, absent in repeated games, even harsher punishments than grim trigger strategies are possible if the marginal cost is positive, namely, to force the defector to price below marginal cost in the post-deviation period, thus bringing down the profit in the deviation period below what it would have been under grim trigger strategies.

In subsection 4.1 we identify a class of such “generalized trigger” strategies, with focus on those with maximally “harsh” punishments, and provide conditions for subgame perfect equilibrium in this type of strategy. Since constant collusive prices that are supported by some subgame perfect strategy profile can be supported by subgame perfect equilibria in such strategies, we thereby explore the full range of stationary subgame perfect equilibrium outcomes, a point to which we substantiate in subsection 4.5. Before that, however, we analyze collusion possibilities in generalized trigger strategies in section 4.2, including limiting results and comparisons with the repeated-games case $\alpha = 1$ and $\beta = 0$. Moreover, in section 4.3 we provide results on the optimal deviation price in the special case of zero marginal cost and monopoly-price collusion, and illustrate the results in section 4.4 by means of a parametric example.

4.1. Generalized trigger-strategy equilibria. Suppose, first, that $c = 0$, and define “grim trigger” strategies just as in the repeated-games case discussed above. Such a strategy profile is a subgame perfect equilibrium if and only if

$$(p - c) A(p^*, p, 0) \leq \frac{\alpha \Pi(p^*)}{n(1 - \delta)} \quad \forall p < p^*, \quad (12)$$

where $c = 0$. The quantity on the left-hand side is the present value of the profit to a firm that undercuts the collusive price by posting a price $p < p^*$ (such a firm will earn zero profit in all later periods), and the quantity on the right-hand side is the present value of the firm’s profit were it to remain at the collusive price p^* — the factor α accounts for the fact that only the young buy in steady state.

Secondly, suppose that $c > 0$. In such cases, harsher punishments than “grim trigger” strategies are in fact possible. The most severe “punishment” of a defector, is to drop the market price as much as possible in the first punishment period, that is below marginal cost all the way down to zero, and to keep the expected continuation profit, as evaluated after a defection, as close as possible to zero. In order to obtain subgame perfection, such severe punishment should be “incentive compatible.” for

the defector to “obey” and for the other firms to implement. We here focus on such strategies, a class of generalized trigger strategies, where the defector prices at zero in the first punishment period, while all other firms price at marginal cost in that period. The defecting firm thus receives all demand in the defection period, and hence makes a loss. In the next period, all firms return to the collusive price p^* for good with probability q , but with probability $1 - q$ the punishment price profile — zero for the defector and marginal cost for the others — is repeated, and so on. The number of punishment periods is thus random. If the defector does not obey the punishment pricing in a punishment period, the others restart the punishment sequence. Other firms have no incentive to deviate from punishing a deviator since they earn zero profit in each punishment period and cannot make a positive profit since at least one other firm prices at or below marginal cost.¹⁴ We assume that the randomization that determines the duration of punishment is public — a heroic but common assumption in the repeated games literature.¹⁵

More precisely, *generalized trigger strategies* are defined as follows. Initially, all firms ask the same price $p^* > c > 0$ and they continue to do so as long as no firm posts a lower price. Suppose a firm i in some period t posts a price $p < p^*$. In the next period, this firm prices at zero while all other firms price at marginal cost: $p_{t+1,i} = 0 < c = p_{t+1,j}$ for all $j \neq i$. This is not a Nash equilibrium in the stage game of that period, however, so incentives have to be created for firms to play along. Let $\varphi : (0, p^*) \rightarrow [0, 1]$, and set $q = \varphi(p)$. With probability q , all firms return to the collusive price p^* in period $t + 2$, while with probability $1 - q$, they keep their prices from period $t + 1$. In the latter case, the same randomization is independently repeated in period $t + 3$, etc., resulting in a geometric distribution of punishment periods (beyond the first). The random number T of punishment periods thus satisfies $\Pr(T \geq 1) = 1$ and

$$\Pr(T = k + 1 \mid T \geq k) = q$$

for positive integers k , where $q = \varphi(p)$. A generalized trigger strategy is thus fully characterized by the randomization function φ .

A generalized trigger strategy profile φ is a subgame perfect equilibrium if and

¹⁴In order to make this a complete strategy specification, assume all firms set the price c also in case current aggregate demand is inconsistent with p^* having been the going and expected price in all preceding periods.

¹⁵An alternative to randomized duration of punishment is to have deterministic non-decreasing pricing schemes during a finite and deterministic number of periods. However, because time is discrete, such strategy profiles lack a certain continuity property that randomized durations have and that allows for certain analytical results.

only if the following two conditions hold for all $p < p^*$ and for $q = \varphi(p)$:

$$(p - c) A(p^*, p, 0) - \delta c B(p^*, p, 0) + \delta^2 \sum_{k=1}^{+\infty} q(1 - q)^{k-1} \left[\left(\frac{1 - \delta^{k-1}}{1 - \delta} \alpha \Pi(0) + \frac{\delta^{k-1}}{1 - \delta} \frac{\alpha}{n} \Pi(p^*) \right) \right] \leq \frac{\alpha \Pi(p^*)}{n(1 - \delta)} \quad (13)$$

and

$$-cB(p^*, p, 0) + \delta \sum_{k=1}^{+\infty} q(1 - q)^{k-1} \left[\left(\frac{1 - \delta^{k-1}}{1 - \delta} \alpha \Pi(0) + \frac{\delta^{k-1}}{1 - \delta} \frac{\alpha}{n} \Pi(p^*) \right) \right] \geq 0. \quad (14)$$

The first term in (13) is the defector's profit in the defection period. Consumers then see the defector's current price p and expect the market price zero in the next period. The second term is the defector's discounted profit from sales in the first punishment period, where $B(p^*, p, 0)$ denotes aggregate demand in this period. The sum on the second line represents the defector's discounted profits during punishment periods 2 and beyond. Only the young buy in these periods, and they make up the population share α . The factor $q(1 - q)^{k-1}$ is the probability that the number T of punishment periods is k . The associated factor in square brackets is the sum of discounted profits when $T = k$; $k - 1$ initial punishment periods of selling to all young at price zero, followed by all firms returning to collusive pricing. The expression on the right-hand side of (13) is the discounted sum of profits that the defecting firm would have earned, had it not defected. Likewise, the first term on the left-hand side of (14) represents the profit to the defecting firm during the first punishment period, and the sum the discounted profits thereafter, evaluated from the first punishment period. The left hand side is thus the present net value to the defecting firm, after its defection, of obeying the punishment. Condition (14) requires this present value to be non-negative; otherwise the defecting firm would do better by pricing at marginal cost forever. A generalized trigger strategy profile that meets condition (14) with equality for all deviation prices $p \in (c, p^*)$ will be called *tight*.

Proposition 1. *Suppose that $\delta \in (0, 1)$, $n \in \mathbb{N}$, $c > 0$ and $p^* \in (c, p^m]$. Condition (12) is necessary for p^* to be supported by subgame perfect generalized trigger strategies. This condition is also sufficient for tight generalized trigger strategies supporting p^* to be subgame perfect. Condition (15) below is sufficient for the existence of tight generalized trigger strategies.*

$$c < \delta \frac{\alpha \Pi(p^*)}{n(1 - \delta)} \quad (15)$$

Proof: Let $\delta \in (0, 1)$, $n \in \mathbb{N}$, $c > 0$, $p^* \in (c, p^m]$ and consider a generalized trigger strategy with randomization function φ , such that p^* is the price on the path of the profile. Then conditions (13) and (14) hold for all $p \in (c, p^*)$ when $q = \varphi(p)$. Insertion of (14) into (13) gives (12), proving the first claim. For the second claim, consider a tight generalized trigger strategy with randomization function φ . Then (14) holds with equality for all $p \in (c, p^*)$. Inserting this equality into condition (13) shows that the latter holds for all $p < p^*$ if and only if condition (12) holds. For the third claim, note, for any given deviation price p , the left-hand side of (14) is an increasing and continuous function f of $q \in [0, 1]$, with $f(0) = -cB(p^*, p, 0)$ and

$$f(1) = -cB(p^*, p, 0) + \delta \frac{\alpha \Pi(p^*)}{n(1-\delta)}.$$

By the intermediate value theorem, $f(0) < 0 < f(1)$ is sufficient for the existence of a termination probability $q \in (0, 1)$ such that (14) holds with equality for that deviation price p . It remains to examine $B(p^*, p, 0)$. This is aggregate demand in the first punishment period, emanating from the young in that period, plus the young in the defection period who did not buy then and survived one period, plus the old in the defection period who did not buy and who survived one more period:

$$\begin{aligned} B(p^*, p, 0) &= \alpha D(0) + (1-\alpha) \alpha F\left(\frac{p}{1-\beta}\right) + \\ &\quad + (1-\alpha)^2 \left[F(p^*) - \max\left\{0, D\left(\frac{p}{1-\beta}\right) - D(p^*)\right\} \right] \\ &= \alpha D(0) + (1-\alpha) \alpha \left[1 - D\left(\frac{p}{1-\beta}\right) \right] + \\ &\quad + (1-\alpha)^2 \left[1 - \max\left\{D(p^*), D\left(\frac{p}{1-\beta}\right)\right\} \right] \end{aligned}$$

Thus $B(p^*, p, 0) > 0$ and hence $f(0) < 0$, and this holds for all p . Moreover, $F(0) \geq 0$ and hence

$$B(p^*, p, 0) \leq \alpha + (1-\alpha)\alpha + (1-\alpha)^2 = 1,$$

again for all p . Thus, condition (15) is sufficient for the existence of a tight generalized trigger strategy profile supporting p^* . In force of claim two, this establishes claim three in the proposition. **End of proof.**

Note that the sufficient condition (15) for the existence of tight generalized trigger strategies is met when $c = 0$, and for all sufficiently large discount factors $\delta < 1$ when $c > 0$. The subsequent analysis is predicated upon condition (15), and is focussed on tight generalized trigger strategies.

4.2. Collusion possibilities: Limit results and comparison with the repeated-games model. Using (6) and dividing through by $\alpha > 0$, the subgame perfection condition (12) can be re-written as

$$(p - c) \left[D \left(\frac{p}{1 - \beta} \right) + \frac{1 - \alpha}{\alpha} \max \left\{ 0, D \left(\frac{p}{1 - \beta} \right) - D(p^*) \right\} \right] \leq \frac{\Pi(p^*)}{n(1 - \delta)} \quad \forall p < p^*. \quad (16)$$

This condition immediately implies a ‘‘Folk Theorem’’- like result:

Proposition 2. *For each collusive price $p^* \in (c, p^m]$ there exists a discount factor $\bar{\delta}(p^*) < 1$ such that p^* is sustainable in subgame perfect equilibrium for all $\delta \in (\bar{\delta}(p^*), 1)$.*

In other words: just as in the standard repeated-games approach, collusive prices are sustainable provided firms are sufficiently patient. Note, however, that unlike in that special case, the critical discount factor in general depends on the collusive price p^* .

We turn to a comparison of the deviation profit in this model with that in the standard repeated-game model, for any given discount factor δ high enough to meet condition (15). The left-hand side of (16) is the present value of the stream of net profits *per young consumer* to a firm that undercuts the going collusive price by setting $p < p^*$ (recall that the young make up the population share α). By continuity of the left-hand side of (16) with respect to p , that inequality holds for all $p < p^*$ if and only if

$$\hat{\pi}(p^*, \alpha, \beta) \leq \frac{\Pi(p^*)}{n(1 - \delta)}, \quad (17)$$

where $\hat{\pi}(p^*, \alpha, \beta)$ is the *maximal deviation profit (per young consumer)*:

$$\hat{\pi}(p^*, \alpha, \beta) = \max_{p \in [0, p^*]} (p - c) \left[D \left(\frac{p}{1 - \beta} \right) + \frac{1 - \alpha}{\alpha} \max \left\{ 0, D \left(\frac{p}{1 - \beta} \right) - D(p^*) \right\} \right]. \quad (18)$$

The right-hand side in the equilibrium condition (17) is identical to that in the standard repeated-games case (inequality (10)). Moreover, if $\alpha = 1$ and $\beta = 0$, as in the standard repeated-games case, then $\hat{\pi}(p^*, \alpha, \beta) = \Pi(p^*)$. Hence, condition (17) then coincides with the repeated-games equilibrium condition (11).

It is easily verified that the generalized deviation profit $\hat{\pi}(p^*, \alpha, \beta)$ is continuous, non-decreasing in p^* , and non-increasing in each of α and β . Hence, collusion is easier the higher the consumer turn-over rate α and their effective discount factor β are, while the effect of a change of a collusive price $p^* < p^m$ is left ambiguous, since both

sides of the inequality (17) are non-decreasing in p^* . However, when $p^* = p^m$, then the right-hand side is maximal, and thus we may conclude that a marginal reduction in the collusive price, from monopoly pricing, does not facilitate collusion. More exactly:

Proposition 3. *The deviation profit $\hat{\pi}(p^*, \alpha, \beta)$ is continuous in all three arguments. For $p^* \leq p^m$, it is non-decreasing in p^* and non-increasing in each of α and β , with $\hat{\pi}(p^*, 1, 0) = \Pi(p^*)$. For each $p^* \in (c, p^m]$ and $\alpha \in (0, 1)$ there exists a $\bar{\beta} \in (0, 1)$ such that*

$$\hat{\pi}(p^*, \alpha, \beta) \geq \Pi(p^*) \quad \Leftrightarrow \quad \beta \leq \bar{\beta}.$$

Proof: The continuity claim follows from the continuity of $D = 1 - F$, by Berge's Maximum Theorem. The equation $\hat{\pi}(p^*, 1, 0) = \Pi(p^*) = \max_{p \in [0, p^*]} \Pi(p)$ follows from the monotonicity of D . Moreover, $\hat{\pi}(p^*, \alpha, \beta)$ is non-decreasing in p^* since the interval $[0, p^*]$, from which the maximand is chosen, is increasing in p^* , and the maximand is point-wise non-decreasing in p^* . Likewise, $\hat{\pi}(p^*, \alpha, \beta)$ is non-increasing in α since the maximand is point-wise non-increasing in α , and likewise for β , see equation (18). Finally, let $p^* > c$ and $\alpha \in (0, 1)$. Then $\hat{\pi}(p^*, \alpha, \beta)$ is continuous and non-increasing in β , with

$$\begin{aligned} \hat{\pi}(p^*, \alpha, 0) &= \max_{p \in [0, p^*]} (p - c) \left[D(p) + \frac{1 - \alpha}{\alpha} [D(p) - D(p^*)] \right] \\ &\geq \max_{p \in [0, p^*]} \Pi(p), \end{aligned}$$

where the inequality follows from the monotonicity of D . **End of proof.**

Recall that $\beta = (1 - \alpha)\gamma$, where γ is the consumers' pure temporal discount factor, and, for now, keep the consumer turn-over rate $\alpha \in (0, 1)$ fixed. It follows from the proposition that the deviation profit is (weakly) lower the more *patient* the consumers are — the higher γ is. The intuition for this is clear. Suppose a firm undercuts the collusive price. The more patient the consumers are, *ceteris paribus*, the more of them will postpone their purchase until next period's anticipated "price war." If consumers are sufficiently *patient*, that is, if $\gamma > \bar{\gamma} = \bar{\beta} / (1 - \alpha)$, then the maximal deviation profit is lower than or equal to that in the repeated-games case: collusion is then easier to sustain than in the repeated-games model. By contrast, if consumers are sufficiently *impatient*, $\gamma \leq \bar{\gamma}$, then the maximal deviation profit (weakly) exceeds that in the corresponding repeated-games case: collusion is then harder to sustain than in the repeated-games model.¹⁶ We illustrate this graphically in an example below.

¹⁶Note that we have not excluded the possibility that $\bar{\gamma} \geq 1$. The analysis under the constraint $\bar{\gamma} \leq 1$ is more cumbersome but does not give much in terms of additional insight.

4.3. Zero marginal cost and monopoly pricing. We here focus on the special case of zero marginal cost, $c = 0$, and collusive monopoly pricing. Unlike in the standard repeated-games case, the optimal deviation is then a sizeable price cut. To see this, first note that

$$\hat{\pi}(p^m, \alpha, \beta) = \max \{ \hat{\pi}_1(p^m, \alpha, \beta), \hat{\pi}_2(p^m, \alpha, \beta) \}, \quad (19)$$

where $\hat{\pi}_1$ is the optimal deviation profit over the price range, $p \leq \bar{p} = (1 - \beta)p^m$, in which both old and young bite, and $\hat{\pi}_2$ is the supremum deviation profit over the (open) price range, $\bar{p} < p < p^m$, in which only young consumers bite.

Proposition 4. $\hat{\pi}_1(p^m, \alpha, \beta) \geq \hat{\pi}_2(p^m, \alpha, \beta)$ for all $\alpha, \beta \in (0, 1)$.

Proof: By way of change of variables (let $r = p / (1 - \beta)$):

$$\begin{aligned} \hat{\pi}_1(p^m, \alpha, \beta) &= \max_{p \leq \bar{p}} \frac{p}{\alpha} \left[D\left(\frac{p}{1 - \beta}\right) - (1 - \alpha) D(p^m) \right] \\ &= \frac{1 - \beta}{\alpha} \max_{r \leq p^m} r [D(r) - (1 - \alpha) D(p^m)] \geq (1 - \beta) \Pi(p^m) \end{aligned}$$

and

$$\begin{aligned} \hat{\pi}_2(p^m, \alpha, \beta) &= \max_{\bar{p} \leq p \leq p^m} p D\left(\frac{p}{1 - \beta}\right) \\ &= (1 - \beta) \max_{r \geq p^m} r D(r) = (1 - \beta) \Pi(p^m). \end{aligned}$$

End of proof.

In other words, it always pays off to make a price all the way down to $(1 - \beta)p^m$ or further, where both old and young bite. Moreover, if D is continuously differentiable, then the optimal undercutting price, \hat{p} , necessarily satisfies the following first-order condition:

$$\Pi' \left(\frac{\hat{p}}{1 - \beta} \right) = (1 - \alpha) D(p^m).$$

In the limit as $\alpha \rightarrow 1$ and $\beta \rightarrow 0$, we obtain $\hat{p} \rightarrow p^m$, just as in the standard repeated-games case.

4.4. Example. For the sake of illustration, suppose that consumer valuations are uniformly distributed on the unit interval: $D(p) = \max\{0, 1 - p\}$. Then $\Pi(p) = \max\{0, p(1 - p)\}$, $p^m = 1/2$, and we obtain

$$\hat{p} = \frac{1}{4}(1 + \alpha)(1 - \beta). \quad (20)$$

Hence, the optimal deviation price increases with the consumer turn-over rate α and decreases with consumers' effective patience β . The optimal deviation profit per young consumer becomes

$$\hat{\pi}_1 = \frac{1 - \beta}{\alpha} \left(\frac{1 + \alpha}{4} \right)^2, \quad (21)$$

a decreasing function both of α and of β . Using $\beta = (1 - \alpha)\gamma$, this combines to

$$\hat{p} = \frac{(1 + \alpha)[1 - (1 - \alpha)\gamma]}{4}, \quad (22)$$

and

$$\hat{\pi}_1 = \frac{1 - (1 - \alpha)\gamma}{\alpha} \left(\frac{1 + \alpha}{4} \right)^2. \quad (23)$$

The diagram below shows isoquants for the optimal deviation price, \hat{p} , as a function of α (horizontal axis) and γ (vertical axis).




Figure 1: Contour map for the optimal undercutting price, as a function of α and γ .

The thick, curve is the isoquant where \hat{p} is just below p^m , that is, where the optimal deviation price is close to the marginal under-cutting that is familiar from the repeated-games model of price competition. From the left, the curves are the isoquants for $\hat{p} = 0.1, 0.2, 0.3, 0.4$ and 0.495 , respectively (recall that $p^m = 0.5$). The optimal deviation price is hence far below the going monopoly price when consumers are patient and their turn-over rate is small, a combination that arises when market periods are short. In other words, the optimal under-cutting is significant when firms cannot commit to their prices for long periods — in stark contrast with the usual repeated-games model of price competition.

The next diagram shows isoquants for the maximal deviation profit, $\hat{\pi}_1$, as a function of α (horizontal axis) and γ (vertical axis). The thick, kinked curve is the isoquant for $\hat{\pi}_1 = \Pi(p^m)$, that is, where the maximal deviation profit is the same as in the repeated-games model. Parameter pairs (α, γ) above (below) this curve result in lower (higher) deviation profits. Hence, collusion is easier (harder) in the present model than in the repeated-games model for parameter pairs above the curve.




Figure 2: Contour map for the maximal deviation profit, as a function of α and γ .

We also note that the two iterated limits of the maximal deviation profit per young consumer, $\hat{\pi}_1$, when $\alpha \rightarrow 0$ and $\gamma \rightarrow 1$, differ:

$$\lim_{\alpha \rightarrow 0} \lim_{\gamma \rightarrow 1} \hat{\pi}_1 = \frac{1}{16} < \Pi(p^m) = \frac{1}{4} < \lim_{\gamma \rightarrow 1} \lim_{\alpha \rightarrow 0} \hat{\pi}_1 = +\infty.$$

However, the iterated limits of the maximal deviation profit *in absolute terms*, $\alpha \hat{\pi}_1$, do agree: both equal $\Pi(p^m) = 1/4$.

To obtain a determined limit of the maximal deviation profit per young consumer, suppose the length of the market period is taken down from one to zero. With a market-period of length $\Delta \in (0, 1)$, we would then have $\alpha = \Delta \cdot a$ and $\gamma = \exp(-b \cdot \Delta)$ for some $a, b > 0$. Taking this continuous-time limit, we obtain

$$\lim_{\Delta \rightarrow 0} \hat{\pi}_1 = \lim_{\Delta \rightarrow 0} \frac{1 - (1 - a\Delta) e^{-b\Delta}}{a\Delta} \left(\frac{1 + a\Delta}{4} \right)^2 = \frac{a + b}{16a}.$$

Hence, in the limit, collusive monopoly pricing is easier (harder) to sustain in the present model than in the repeated-games model if $b < 3a$ ($b > 3a$), that is, if consumers are relatively patient in comparison with their turn-over rate.

Finally, we note that collusion at the monopoly price constitutes a *strict* equilibrium for many parameter combinations, in the sense that the maximal deviation profit is strictly lower than the equilibrium profit. This is, for example, the case for all (α, γ) -pairs below the thick curve in Figure 2, granted condition (11) holds.

4.5. Other strategies. So far we have focused on a certain class of strategies — tight generalized trigger strategies. The reader might wonder to what extent our results are predicated on this restriction. We believe they are not and in this section we will discuss why. In particular, we claim that any collusive price that can be sustained in subgame perfect equilibrium can also be sustained in subgame perfect equilibrium when firms use tight generalized trigger strategies — a conclusion that is a generalization of that for “grim” trigger strategies in the standard repeated-game model. The reasoning in the current model has one twist that is absent in the repeated game. The conclusion holds in the repeated-game case because *future* profits — after a deviation — are as low as possible (zero) in the subsequent continuation subgame when grim trigger strategies are used. The same holds true in the present model when $c = 0$, and then also *current* deviation profits are as low as possible, because consumers have the greatest incentive to postpone purchases if they anticipate the lowest possible price (zero) the next period. For $c > 0$, however, the latter no longer holds true for grim trigger strategies: *current* period deviation profits are instead lowest under tight generalized trigger strategies, because consumers then anticipate the price zero, instead of $c > 0$ in the next period.

To see this, consider any subgame perfect strategy profile supporting a constant collusive price p^* . The implied punishment is then necessarily milder than under tight generalized trigger strategies. A consumer with perfect foresight anticipates the path of prices $p^e(\tau)$, $\tau = 1, 2, \dots$ under the punishment strategy profile in question, where the (current) deviation period has been labelled $\tau = 0$. The consumer thus buys in the current period if and only if condition (2) holds. It immediately follows that the quantity on the right-hand side is not larger than under tight generalized trigger

punishment — since then $p^e(1) = 0$. Hence, the incentives of consumers to wait until the next period is the highest under tight generalized trigger strategies. The current period deviation profit is thus (weakly) larger for a deviating firm under any other strategy profile. This is a new phenomenon that does not exist in the standard repeated-game model: intertemporal consumers are more likely to buy in the current period if they anticipate a milder price war. The second effect — which does exist in the standard model — is that a milder punishment also implies that future profits are higher, as compared with those under tight generalized trigger strategies. Therefore, the two effects work in the same direction and together imply that (total) deviation profits are higher under any alternative strategy profile. In sum:

Proposition 5. *Any collusive price that can be sustained in subgame perfect equilibrium can be sustained in subgame perfect equilibrium under tight generalized trigger strategies.*

5. EQUILIBRIUM SALES

We have so far only considered constant collusive prices. However, a temporary price cut, in equilibrium, could be a way for the industry to capture residual demand and thereby increase firms' joint profits above the monopoly profit for the industry. Temporary equilibrium price cuts are qualitatively different than out-of-equilibrium price cuts, since in the first case, consumers do not expect an ensuing price war and hence do “bite” during the sale. A model of sales has been developed in Sobel (1984, 1991). However, while no consumer “dies” in Sobel’s models, and therefore residual demand under a constant market price builds up without bounds over time, our consumers do “die,” and residual demand thus is bounded. Nevertheless, residual demand may be sufficiently large to motivate temporary sales in equilibrium. We here identify conditions under which “sales equilibria” with profits above monopoly profits do and do not exist. Roughly speaking, existence hinges upon whether or not firms are more patient than consumers — an observation qualitatively in line with Sobel’s findings. Similar results to the ones presented here have been obtained by Argenton (2004) for the case of two-point valuation distributions.

In order to highlight the potential profitability of equilibrium sales, let us first briefly consider an extreme case. Suppose that consumers are maximally impatient, $\gamma = 0$, and that all firms ask the monopoly price, p^m , in all “normal” periods, and some lower price p^s in a unique sales period, $t = S$. Then each firm earns its share of the industry monopoly revenue, $\alpha\Pi(p^m)/n$, in all normal periods, just as under constant collusion at the monopoly price. For although new consumers anticipate the upcoming sale, their impatience drives them to buy in their first period in the market. In the sales period, however, also all old consumers with valuations between

p^s and p^m buy. Choosing the sales price p^s optimally, all firms earn more in that period than at the monopoly price:

$$\begin{aligned} & \max_{p^s \in [0, p^m]} \frac{p^s - c}{n} [\alpha D(p^s) + (1 - \alpha) \max \{0, D(p^s) - D(p^m)\}] \\ & \geq \frac{\alpha}{n} \max_{p \in [0, p^m]} \Pi(p) = \frac{\alpha}{n} \Pi(p^m). \end{aligned} \quad (24)$$

Generically, this inequality holds strictly, and, by continuity, it then also holds for all $\gamma > 0$ sufficiently close to zero. Moreover, the deviation incentive, both in “normal” and “sales” periods, can be analyzed along the lines developed above for steady-state collusion.

More generally and precisely, consider recurrent sales, where all firms set the same sales price p^s every S periods, where $S > 1$ is an integer, and ask the same price $p^o \geq p^s$ in all other, “ordinary,” periods. In all non-sales periods, only new consumers buy, and only those with valuations above $v^\sigma(p^o, p^s)$, where σ is the number of periods remaining before the next sale, see equation (3). All other young consumers in the cohort will postpone purchase until the next sale, or not buy at all. Hence, if the sale occurs in periods $0, S, 2S, 3S, \dots$, then industry profits (per new consumer) in periods $t = 1, 2, \dots, S - 1$ are

$$\Pi_t(p^s, p^o) = (p^o - c) (1 - F[v^{S-t}(p^o, p^s)]),$$

and

$$\Pi_S(p^s, p^o) = (p^s - c) [1 - F(p^s)] + \sum_{t=1}^{S-1} (1 - \alpha)^{S-t} (p^s - c) (F[v^{S-t}(p^o, p^s)] - F(p^s)).$$

We note that the critical valuation for purchase in a pre-sales period, $t < S$, is creasing in t :

$$v^{S-t}(p^o, p^s) = \frac{p^o - \beta^{S-t} p^s}{1 - \beta^{S-t}} = p^o + \frac{p^o - p^s}{\beta^{t-S} - 1}.$$

Consequently, aggregate demand and industry profits decline over time t as the sale in period S comes closer. The present value of the expected stream of future profits thus varies over each sales cycle.

We report two results, one positive and one negative, for the profitability and sustainability of sales equilibria of this form. For the positive result, suppose parameters are such that p^m is sustainable as a strict steady-state subgame perfect equilibrium (an example was given in section 4). Does there then exists a subgame perfect strategy profile, with recurrent sales as defined above, which results in a higher discounted

payoff to all firms? In order to answer this question, we consider the discounted sum of industry profits during a whole cycle, as evaluated from the first period after a sale:¹⁷

$$\Psi(p^s, p^o) = \sum_{t=1}^S \delta^t \Pi_t(p^s, p^o).$$

For an affirmative answer to the posed question, it is sufficient to show that there exists a price pair (p^s, p^m) such that (a) $\Psi(p^s, p^m) > \Psi(p^m, p^m)$ and (b) there exists a subgame-perfect strategy profile that induces the sales price-pair (p^s, p^m) .

Proposition 6. *Suppose that the valuation distribution F is differentiable with positive density at the monopoly price, p^m , and that the monopoly price is sustainable as a strict steady-state subgame perfect equilibrium. If $\delta > \gamma$, then there exists a continuum of subgame perfect sales price-pairs (p^s, p^m) that yield higher discounted payoffs to all firms than collusion at the monopoly price.*

Proof: We first establish that $\Psi(p^s, p^m) > \Psi(p^m, p^m)$ for all $p^s < p^m$ sufficiently close to p^m by way of taking the derivative of $\Psi(p^s, p^m)$ with respect to its first argument, at $p^s = p^m$:

$$\begin{aligned} \Psi_1(p^s, p^m) \Big|_{p^s=p^m} &= \sum_{t=1}^{S-1} \delta^t \frac{(p^m - c) F'(p^m)}{\beta^{t-S} - 1} \\ &+ \delta^S \left[\Pi'(p^m) - \sum_{t=1}^{S-1} (1 - \alpha)^{S-t} \frac{\beta^{t-S} (p^m - c) F'(p^m)}{\beta^{t-S} - 1} \right] \\ &= \delta^S \sum_{t=1}^{S-1} \frac{\delta^{t-S} - (1 - \alpha)^{S-t} \beta^{t-S}}{\beta^{t-S} - 1} (p^m - c) F'(p^m), \end{aligned}$$

where we used the first-order condition $\Pi'(p^m) = 0$. By hypothesis, $F'(p^m) > 0$ and $p^m > c$. Hence, $\Psi_1 < 0$, if $\delta^{t-S} < (1 - \alpha)^{S-t} \beta^{t-S}$. Using $\beta = (1 - \alpha)\gamma$, the last inequality is equivalent with $\delta > \gamma$. This establishes our first claim. The second claim follows by continuity: since by hypothesis p^m is a strict equilibrium price, generalized trigger strategies will also support, in subgame perfect equilibrium, sales prices $p^s < p^m$ sufficiently close to p^m . **End of proof.**

In other words, if firms are more patient than consumers, then there exist sales equilibria that result in profits above monopoly profits. What if consumers are more

¹⁷The following results do not depend on this choice: they can be shown to be independent of which period in the cycle is taken as the point of evaluation.

patient than the firms? At least in the special case of linear demand, it is not difficult to show that sales equilibria are unprofitable:

Proposition 7. *Suppose that the valuation distribution F is uniform and $c < p^* \leq p^m$. If $\delta \leq \gamma$, then $\Psi(p^s, p^*) < \Psi(p^*, p^*)$ for all $p^s < p^*$.*

Proof: Let $p^* \in (c, p^m]$ be given, and take the derivative of $\Psi(p^s, p^*)$ with respect to its first argument, at any $p^s \leq p^*$:

$$\Psi_1(p^s, p^*) = \delta^S \left[\Pi'(p^s) + \lambda \sum_{t=1}^{S-1} \frac{\delta^{t-S} - (1-\alpha)^{S-t} \beta^{t-S}}{\beta^{t-S} - 1} (p^* - c) \right],$$

where $\lambda > 0$ is the constant density of F . We have $\Pi'(p^s) > 0$ for all $p^s < p^m$, and each term in the sum is non-negative if $\delta^{t-S} \geq (1-\alpha)^{S-t} \beta^{t-S}$, or, equivalently, if $\delta \leq \gamma$. **End of proof.**

6. CONSUMERS WITH ADAPTIVE EXPECTATIONS

Some generalizations to consumers with imperfect foresight seem analytically feasible in our modelling framework. For example, if all consumers have adaptive expectations in the sense that they always expect the current price to prevail also in the future, then a deviating firm will sell to all new and old consumers with valuations above the under-cutting price, and hence earn a higher profit than when consumers have perfect foresight. Consequently, it is harder to collude against consumers with adaptive expectations.

To see the implications of adaptive expectations more precisely, note that such consumers' behavior is identical with that of consumers with perfect foresight but with maximal impatience. The definition of aggregate demand in equation (6) is valid for any expectations, not just perfect foresight, and all of the above analysis applies to adaptive expectations, by way of either setting consumers' price expectation equal to the current price, $p^e = p$, or, equivalently, by assuming perfect foresight but setting the effective discount factor β equal to zero. It follows immediately from Proposition 3 that collusion against consumers with adaptive expectations ($\beta = 0$) is easier than collusion against consumers with rational expectations ($\beta > 0$). It also follows that collusion against consumers with adaptive expectations and intertemporal substitution possibilities ($\beta = 0$ and $\alpha < 1$) is easier than collusion in the repeated-games model ($\beta = 0$ and $\alpha = 1$), see Proposition 3 and the example in section 4.4.

The current model also allows for intermediate cases, when either a population fraction $\lambda \in (0, 1)$ of the consumers have perfect foresight and the rest adaptive

expectations, or where all consumers form expectations that are a convex combination of adaptation and perfect foresight. For instance, if $p < p^*$ is the current price and firms use generalized trigger strategies, then adaptive expectations give $p^e = p$, perfect foresight $p^e = 0$, and intermediate expectations $p^e = \mu p$ for some $\mu \in (0, 1)$.

7. CONCLUDING COMMENTS

Our model of Bertrand competition in recurrent market interaction is built on a number of simplifying assumptions. One such assumption is that consumers are homogeneous with respect to their time preferences: all consumers are assumed to have the same discount factor β . A less restrictive assumption would be to assume that β , like the valuation v of the good, is drawn from some fixed probability distribution.

Another simplification is that we have focused on the case of an indivisible good. It seems likely that the qualitative results carry over also to the case of divisible goods. A third simplification is the assumption of no resale. For many durable goods, there are second-hand markets, and these markets interact in an important way with the markets for new units. These are tasks for future research.

Also, we have focused exclusively on the idealized case of perfect foresight on behalf of the consumers. Some generalizations to consumers with imperfect foresight seem analytically feasible in our modelling framework, see comments in the preceding section. A study of collusion against boundedly rational consumers would be a valuable extension.

We finally mention yet two other avenues for future research, namely to apply the present model of intertemporal demand to Cournot oligopoly and to oligopolistic competition when firm's products are imperfect substitutes.

REFERENCES

- [1] Argenton, C. (2004): "Equilibrium sales in Bertrand oligopoly with intertemporal demand", Boston University, mimeo.
- [2] Ausubel, L., and R. Deneckere (1987): "One is almost enough for monopoly", *Rand Journal of Economics* 18, 255-274.
- [3] Coase, R. (1972): "Durability and monopoly", *Journal of Law and Economics* 15, 143-149.
- [4] Conlisk, J., E. Gerstner, and J. Sobel (1984): "Cyclic Pricing by a Durable Goods Monopolist", *Quarterly Journal of Economics* 99, 489-505.
- [5] Dutta, P. (1995): "A Folk theorem for stochastic games", *Journal of Economic Theory* 66, 1-32.

- [6] Fudenberg, D., and J. Tirole (1991): *Game Theory*, MIT Press, Cambridge, MA.
- [7] Gul, F., H. Sonnenschein, and R. Wilson (1986): "Foundations of dynamic monopoly and the Coase conjecture", *Journal of Economic Theory* 39, 155-190.
- [8] Gul, F. (1987): "Noncooperative collusion in durable goods oligopoly", *Rand Journal of Economics* 18, 248-254.
- [9] Maskin, E., and J. Tirole (1988): "A theory of dynamic oligopoly, II: Price competition, kinked demand curves and fixed costs", *Econometrica* 56, 571-600.
- [10] Kirman, A., and M. Sobel (1974): "Dynamic oligopoly with inventories", *Econometrica* 42, 279-287.
- [11] Radner, R. (1999): "Viscous demand", mimeo., Stern School, NYU, New York.
- [12] Selten, R. (1965a): "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträgeit - Teil I: Bestimmung des dynamischen Preisgleichgewichts", *Zeitschrift für die gesamte Staatswissenschaft* 121, 301-324.
- [13] Selten, R. (1965b): "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträgeit - Teil II: Eigenschaften des dynamischen Preisgleichgewichts", *Zeitschrift für die gesamte Staatswissenschaft* 121, 667-689.
- [14] Sobel, J. (1984): "The Timing of Sales", *Review of Economic Studies* 51, 353-368.
- [15] Sobel, J. (1991): "Durable Goods Monopoly with Entry of New Consumers", *Econometrica* 59, 1455-1485.
- [16] Tirole, J. (1988): *The Theory of Industrial Organization*, MIT Press, Cambridge, MA.
- [17] Wallner, K. (1999): "Sequential moves and tacit collusion: Reaction function cycles in a finite pricing duopoly", *Journal of Economic Theory*, 84, 251-267.