BERTRAND COMPETITION WITH INTERTEMPORAL DEMAND

Prajit Dutta
Department of Economics
Columbia University

Alexander Matros
ELSE, Department of Economics
University College London

Jörgen W. Weibull*
Department of Econmics
Boston University

20 March 2003

Abstract. The text-book model of dynamic oligopolistic competition views firms as players in a repeated game, where the demand function is the same in every period. We argue that 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 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 recurrent market setting, and use it to analyze collusive pricing under Bertrand competition. The more patient and forward-looking consumers are, the easier it is for firms to collude against them.

*This is a revised version of WP 493 at the Stockholm School of Economics. The authors are grateful for helpful comments from Cedric Argenton, Avinash Dixit, Roman Inderst, Sergei Izmalkov, seminar participants at Boston University, Columbia University, and University College London, as well as participants at the Annual Congress of the European Economic Association 2001 and at the First Annual International Conference, "Russia 2015: A Long-Term Strategy", 2001. Matros and Weibull thank the Hedelius and Wallander Foundation for financial support. Most of Weibull's research was carried out at the Stockholm School of Economics and at the Research Institute of Industrial Economics, Stockholm.

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 this price has been announced, the monopolist will have an incentive to reduce his 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 also 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, as the literature on repeated games suggests they may, then consumers might foresee such 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 analyze 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.[2]

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,

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

[2]However, 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 - in our recurrent market setting.

see for example Gul, Sonnenschein and Wilson (1986), Gul (1987) and Ausubel and Deneckere (1987). 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 the market in each market period. In this respect we follow Sobel (1984, 1991). However, while the consumers in his model are infinitely lived, and hence residual demand builds up indefinitely over time, our consumers have a finite (random) life time – some "old" consumers leave 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.

More precisely, there is a continuum population of consumers. The market is open over an infinite sequence of market periods, each of length $\lambda \in [0, 1]$. In each market period every firm commits to an ask price. A new cohort of consumers enters the market in each market period, and an equally large set of "old" consumers – that is, who lived in the preceding period – leave the market. The size of the consumer population is thus constant. The population share of new consumers in any period is some fixed positive fraction, which we normalize to λ; this is also the share of the previous period's population that "dies." All consumers have the same probability of dying each period. The life span of every individual is thus a geometrically distributed random variable with constant hazard rate λ (the probability of dying in the next period, when alive in the current period) and with expected value $1/\lambda$ time periods, that is, 1 time unit. The demographic composition of the (continuum) population of consumers is thus treated as deterministic and stationary, while the life span of every individual is random.

The good in question is assumed to be sold in indivisible units, and each consumer is assumed to want to buy the good at most once during his or her lifetime. Once purchased, the good gives a certain utility to the consumer. While having identical life table distributions, consumers differ as to their utility 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.[3] 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-

[3]Note that we do *not* say that consumers who buy a unit "leave" the population. They remain in the population of consumers until they "die", but are no longer active in the market.

theoretic sense but model consumers as price-taking and expectation-forming economic agents with no strategic incentive or power. Their aggregate will constitute a state variable in a stochastic game played by the firms. 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 more or less imperfect foresight. Our position is rather that the contrast in current 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 and sophistication 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 less familiar class of stochastic games (which contains repeated games as a subclass). We here outline how such a generalization can be made, and what are its most direct implications. We also believe that our model of demand can be a useful work-horse for other dynamic market analyses, and for analyses of consumers with imperfect foresight.

The paper delivers one key trade-off and two 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) price-cut. In the standard model, a price cut induces consumers to simply switch their buying *en masse* to the firm that under-cuts the going price. When consumers are forward-looking, however, 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 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 old consumers who have not yet purchased the good; this increases the attractiveness of deviating. The relative size of these two countervailing forces determines the payoff to deviation.

We compare that payoff to the analogous deviation payoff in the textbook model. The parameter of greatest importance turns out to be consumer patience. When consumers are patient, the collusive force prevails; (most) consumers, young and old, prefer to wait for the price war, thereby lowering the deviation payoff to undercutting, and thereby discouraging deviations from the collusive price in the first place. On the other hand, when consumers are impatient, the fact that there are old customers who add their demand to that of the newly born ones imply that deviation payoffs are higher than in the textbook model. Hence collusion is then harder to sustain than in the textbook model.

One other interesting conclusion emerges from the analysis and that is the importance of "sales." In the textbook model, if a firm were to undercut a collusive price, say p^*, then it would do so by undercutting ever so marginally. The reason for that is that any price below p^* would guarantee the whole market for the deviant, and the whole market's revenue is highest at p^* (in the standard case when p^* does not exceed the monopoly price and the industry revenue function is single-peaked). In the current model, by contrast, the most profitable under-cutting price is substantially less than the collusive price. This is done in order to attract old consumers. Only a significant price cut can bring in old consumers, and since they represent a positive fraction of the potential buyers, a deviant firm, when the collusive price is at or near the monopoly price, would find it worthwhile to capture their demand.

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 (1965) 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.[4] Closer to our model, however, are the models in Sobel (1984, 1991), where a new cohort of consumers enters the market each period. However, in those models all consumers stay in the market until they make a purchase. Since there is no other "exit" for consumers, this implies that residual demand grows over time when the market price is held constant. This explains Sobel's finding that firm(s) now and then reap the accumulated residual demand by a temporal reduction in the ask price. The price thus stays high for some time, then drops to a low level, and then returns to the high level etc. To the best of our knowledge, we are the first to model oligopolistic competition with demand emanating from a demographically stationary population of intertemporally substituting and forward-looking consumers.

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, concerning subgame-perfect equilibrium in grim trigger strategies, are given in section 4, and illustrated by means of an example in section 5. Section 6 discusses general subgame-perfect

[4] For a discussion of such games and an equilibrium characterization, see Dutta (1995).

equilibria, and section 7 concludes.

2. The Model

2.1. Firms. Suppose there are n firms in a market for a homogenous indivisible good. The market operates over an infinite sequence of market periods, $t \in \mathbb{N} = \{0, 1, 2, ...\}$, each of length $\lambda > 0$. All firms simultaneously announce their *ask prices* at the beginning of every period and are committed to that price during the period. Let $p_{it} \geq 0$ be the price that firm i asks in period t, and let $p_t = (p_{1t}, ..., p_{nt})$ be the vector of ask prices in that period. All consumers observe all current ask prices, and when they buy, they buy only from the firms with the lowest current ask price. The lowest ask price in any period t,

$$p_t = \min\{p_{1t}, ..., p_{nt}\}, \tag{1}$$

will accordingly be called the *market price* in that period. If more than one firm asks the lowest price, then sales are split equally between these. The firms face no capacity constraint and incur no production costs. Hence, each firm's profit in a market period is simply its sales multiplied by its ask price. All firms discount future profits by the same *discount factor* $\delta \in (0, 1)$ between successive market periods.

We assume that firms (a) have complete information about the market structure and the law of motion of consumer demand (see next subsection), (b) set their ask prices simultaneously in each period, and (c) know all past ask prices. A (pure) *strategy* for a firm is thus a function or rule that specifies its ask price in each period, given any history of ask prices in all earlier periods. Firms' strategies constitute a *subgame perfect equilibrium* if in all periods each firm maximizes its expected discounted future stream of profits conditional on any price history before that period, and given all other firms' strategies.

2.2. Consumer demographics. There is a continuum population of consumers. A cohort of new consumers are "born" at the beginning of each period, and an equally large set of "old" consumers – that is, who lived in the preceding period – "die" at the beginning of the period.[5] Consumers who already made a purchase and for that

[5]The analysis does of course not require that "birth" and "death" are interpreted literally. What matters is that these events mark entry and exit of individuals in this market (recall, however, that we do not call a purchasing consumer an exiting consumer). One may, for example, suppose that individuals live for a long time but participate in this market only during a random stretch of consecutive periods.

reason became inactive in the market are still included in the consumer population, and will sometimes be referred to as *satisfied consumers*. The size of the consumer population is thus constant, and we normalize it to 1. We model the demographic process over the sequence of recurrent market periods as the result of an underlying constant flow of births and deaths in continuous time. Hence, the population share of new consumers in any period is $\lambda \in (0,1]$, the length of a market period, and this is also the share of the previous period's population that dies at the beginning of the period in question.

Furthermore, all consumers have the same probability of dying each period. Each consumer who is alive in any period has probability $1-\lambda$ of being alive next period as well. The life span S of any individual is thus a geometrically distributed random variable, with probability λ for $S=\lambda$, probability $(1-\lambda)\lambda$ for $S=2\lambda$, etc. The expected life span of an individual is thus $E[S]=1$.[6] The demographic composition of the population of consumers is thus deterministic and stationary, while the life span of every individual is random, with the same probability $1-\lambda$ of surviving from one period to the next. Since the survival probability is independent of age, it will be sufficient to distinguish between *new* consumers, that is, those who were born in the current market period, and *old* consumers, that is, all the others (who were born in some earlier market period). We initialize this demographic process by assuming that the population shares of new and old consumers in the initial market period, $t=0$, are λ and $1-\lambda$, respectively. Hence, these are the population shares in *all* periods.

Each consumer wishes to buy at most one unit of the good during his or her lifetime. While having identical life table distributions (and time preferences, see below), consumers differ as to their valuation v of the good. In each period, the newborn individuals' valuation v is distributed according to some cumulative distribution function $F:\mathbb{R}_+ \rightarrow [0,1]$. The expected life-time utility to an individual with valuation v from buying a unit of the good τ periods later (for $\tau \geq 0$) at price p is $\beta^\tau (v-p)$, where $\beta \in [0,1)$ is every consumer's discount factor, representing consumers' time preference. The life-time utility of never buying a unit of the good is normalized to zero. Hence, v is the consumer's *maximal willingness to pay*, in

[6]In steady state, the share of newly born is thus λ, the share of one-period olds is $(1-\lambda)\lambda$, of two-periods olds is $(1-\lambda)^2\lambda$, and - more generally - of T periods olds is $(1-\lambda)^T\lambda$, for $T=0,1,2,.....$ Even outside of steady-state, the population share of newly born is always λ, but the distribution of other ages then depends on initial population shares and calendar time.

the following non-strategic sense: a consumer, new or old, who does not yet own the good and who has valuation v, derives (life-time) utility $v - p$ from buying one unit of the good at price p in the current period, to be compared with the (life-time) utility 0 of never buying the good. A consumer who has bought a unit becomes a *satisfied consumer* (but remains, for accounting reasons, in the consumer population until drawn for a random exit, as described above).

We assume the c.d.f. F to be continuous, and define $D(p) = 1 - F(p)$ as the *static aggregate demand* at (market) price p; this is the quantity that is sold when $\lambda = 1$, that is, when the whole population consists of new consumers. We also assume that resale is not possible, and that consumers hold identical expectations about future prices. The main focus of the subsequent analysis is on intermediate consumers time preferences, that is, $\beta \in (0, 1)$, but we will also pay attention to the following two limiting cases: when consumers are *minimally patient*, $\beta = 0$, and when they are *maximally patient*, $\beta \to 1$.[7]

2.3. Consumer expectations and consumer choice. In the case when the consumer population does not turn over every period, that is, when $\lambda < 1$, there are two distinct consumer groups in each market period: new and old. Aggregate demand in any market period can accordingly be decomposed into demand from new consumers, and residual demand from those old individuals who have not yet bought a unit of the good. For both subpopulations, current demand depends on the current market price as well as on current expectations about future prices. Additionally, the demand of the old also depends on the distribution of valuations for those amongst them who have not yet made a purchase. We here discuss these two new elements of the model: consumer expectations of future prices and the distribution of residual valuations.

Consumers, both young and old, are assumed to have point expectations about prices in future periods. Moreover their expectations will be assumed to be identical, that is, all consumers, young and old, will have the same expectation of the market price in any given future period. This expectation can - and typically will be - based on the current market price p, and will be denote $p^e(p, \tau)$ for the market price τ periods from today.[8] Of particular relevance for the subsequent analysis is the

[7]Note, however, that if the good is durable, and if consumers actually die when leaving the market, and if they have no bequest motive concerning their unit of the good, then it would be natural to impose $\bar{\beta} = 1 - \lambda < 1$ as an upper bound on β.

[8]Note that a more general formulation would be one in which the price expectation is formed on

special case when the future market price is expected to remain constant, that is, when $p^e(p,\tau) = p^e(p)$ for all $\tau > 0$. Sometimes – in order to save on notation – we will suppress the current market price and write the expectation of a constant future price simply as p^e.

If an individual with valuation v, who has not yet bought a unit, faces a current market price p and expects the market price τ periods later to be $p^e(p,\tau)$, for $\tau = 1, 2, ...$, then it is optimal to buy in the present period if and only if her expected utility from doing so is neither exceeded by the utility from never buying the good, nor by the expected utility from buying the good in some future period, that is, iff

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

When the future market price is expected to remain constant, that is, when $p^e(p,\tau) = p^e(p)$ for all $\tau > 0$ the analysis specializes in the following ways. First, suppose that the expected future price equals the current price: $p^e(p) = p$. Then all consumers with valuations $v \geq p$, and who have not yet bought a unit of the good, will buy a unit in the current period, while individuals with lower valuations find the good too expensive to be worth-while buying at all.[9] Secondly, and this is a relevant scenario in the subsequent analysis, suppose that the expected future price is zero. After the current period, the best period to buy is then clearly the next one, since $\beta < 1$. The utility from doing so is $\beta v \geq 0$. Hence, it is optimal to buy today iff $v - p \geq \beta v$. Thus, precisely those consumers who have not yet bought the good and who have valuations $v \geq p/(1-\beta)$ will buy in the present period.

Third, and more generally, the *cut-off valuation* level for purchase, when the current price is p, and all future expected prices are p^e, is

$$v^+(p, p^e) = \frac{p - \beta p^e}{1 - \beta}, \tag{3}$$

granted $\beta p^e \leq p$; otherwise it is optimal not to buy at all.

---

the basis of all prices observed in the past, and where it may also depend on calendar time. Since different cohorts of consumers have different lengths of histories to condition on, this introduces the possibility that different generations might have different expectations about future prices. In turn that contradicts the hypothesis of perfect foresight, a hypothesis that we will work with in the subsequent analysis. To avoid that contradiction we make the simpler assumption that expectations are only based on information that is commonly observed, i.e., the current price.

[9]For the sake of definiteness, but without affecting the results, we assume that consumers who are indifferent between buying now and in the future or not at all, will buy now.

3. Preliminaries

Having identified the cut-off valuation for all consumers, we next consider the distribution of residual valuations, that is, the valuations of those old consumers who have not bought the good in an earlier period. This residual valuation distribution will determine residual demand.

3.1. Distribution of residual valuations. Consider a fixed cohort, say, consumers who were all born τ periods in the past. From the analysis of equations (2) and (3) it is clear that it is always the upper tail of the value distribution that buys. Hence, τ periods back, everybody above some cut-off valuation purchased the good. In the next period, that is, $\tau - 1$ periods back, some more consumers from the same cohort may have purchased the good at a lower price. These additional buyers belong to the upper tail among the non-buyers from τ periods back. And so on. So by the time this cohort reaches the current period there is a cut-off valuation $v^-(\tau)$ such that all consumers with valuation above $v^-(\tau)$ have already purchased the good and none of those with valuation below $v^-(\tau)$ have done so.

It follows that the value distribution among the old consumers in any given period who have not yet made a purchase, is completely described by the vector of past cut-off valuations $v^- = (v^-(1), v^-(2), ...v^-(\tau), ...)$. In the special case of a constant past price p^* and constant past price expectation $p^e(p^*, \tau)$, these cut-off valuations are the same, that is, $v^-(\tau)$ is identical for all past cohorts τ.

The vector of past cut-off valuations v^- determines the aggregate demand emanating from all old consumers. We turn to that next.

3.2. Aggregate demand. Consider a period with market price p and a constant expected future price p^e. The aggregate demand emanating from new consumers (those born in the current period) is simply their population share, λ, times the number of new consumers with valuations exceeding the current cut-off valuation, defined in equation (3):

$$\lambda \left(1 - F\left[v^+\left(p, p^e\right)\right]\right). \tag{4}$$

The aggregate demand emanating from the old, by contrast, also depends on past cut-off valuations. For instance, if the market price in all past periods was constant, say p^*, and had always been expected to remain constant at that level, then all individuals with higher valuations would have already bought a unit, and those with lower valuations would have abstained from buying, that is, we would then have $v^-(\tau) = p^*$ for all past cohorts $\tau = 1, 2, ..$ (this can be verified in equation

(3)). In particular, if $p \geq p^*$, there would be no current demand amongst the old, while if $p < p^*$, some additional demand would emanate from the old. Hence, with all past prices and price expectations equal to p^*, aggregate current demand from the old, whose population share is $1 - \lambda$, would be

$$(1-\lambda)\max\left\{0, F(p^*) - F\left[v^+(p, p^e)\right]\right\}. \tag{5}$$

Current total demand, the sum of (4) and (5), thus depends on three prices - the earlier price (and earlier price expectation) p^*, the current price p, and the currently expected future price, p^e:

$$D(p^*, p, p^e) = \lambda D\left[v^+(p, p^e)\right] + (1-\lambda)\max\left\{0, D\left[v^+(p, p^e)\right] - D(p^*)\right\}. \tag{6}$$

To summarize, consumers buy in the current period if they have not bought yet and their valuation is above a certain cut-off point $v^+(p, p^e)$. In turn, that cut-off value forms part of the state variable in the next period – the state in any period being fully determined by the vector v^- of residual valuations. Aggregate demand, and hence firms' profits, are determined from this state variable and the actions taken by the firms, their current ask prices, through equation (6). The strategic interaction between the firms over time is therefore not a repeated but a stochastic game.[10]

3.3. Trigger-strategy equilibrium. Despite this not being a repeated game, *trigger strategies* can be defined in much the same way as in a repeated game: All firms ask the same price p^* in the initial period ($t = 0$), and continue to do so in all future periods as long as all firms quote that price. If any firm asks any other price in a period, then all firms ask the price 0, their marginal production cost, in all subsequent periods.[11]

We assume that consumers have perfect foresight in the sense that they correctly anticipate pricing under trigger strategies, at least on the price path induced by such strategies and after any unilateral deviation,

$$p^e(p) = \begin{cases} p^* & if\ p = p^* \\ 0 & if\ p \neq p^* \end{cases}, \tag{7}$$

[10]For 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.

[11]Recall that the marginal cost is zero, so for all firms to quote the price zero, irrespective of the current population state, is a Nash equilibrium of the stage game in any given period.

where p is the current market price, as defined in equation (1). In turn, firms correctly compute residual valuations, and hence also aggregate demand, when they take this price expectation function into account.

We shall now investigate the sustainability of trigger-strategy profiles in this recurrent market setup. When testing for subgame perfection of such a strategy profile, for a given "collusive" price p^*, we will assume that the initial demographic state is consistent with p^* having been asked in the past.[12] Other strategies are considered in Section 6.[13]

3.4. The textbook model. The standard analysis of dynamic oligopoly is clearly a special case of this recurrent market setup. More precisely, when consumers live for and care about one period only, that is, when $\lambda = 1$ and $\beta = 0$, then the present model is nothing but the standard repeated-games model of Bertrand competition: the oligopoly then faces the same demand function $D = 1 - F$ each period. We will refer to this boundary case as *the textbook model* and use it as a benchmark for the subsequent analysis.

When the n firms face the same demand function D in each period, then 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 zero, constitutes a subgame-perfect equilibrium if and only if

$$pD(p) \leq \frac{p^* D(p^*)}{n\,(1-\delta)} \tag{8}$$

for all $p < p^*$. The quantity on the left-hand side is the profit to a firm which undercuts the collusive price by asking a lower price p, and the quantity on the right-hand side is the present value of each firm's profit at the collusive price p^*. By continuity of the demand function, inequality (8) holds for all $p < p^*$ if and only if

$$\hat{\pi}^t\,(p^*) - \pi^c\,(p^*, n, \delta) \leq 0, \tag{9}$$

where

$$\hat{\pi}^t\,(p^*) = \sup_{p<p^*} pD(p) = \max_{p\in[0,p^*]} pD(p) \tag{10}$$

[12] More precisely, we assume that in period $t = 0$, all old consumers with reservation values $v \geq p^*$ already own a unit, so that the valuations of the remaining old are given by the c.d.f. F truncated at p^*.

[13] In particular, when we write that such and such an outcome is sustainable in subgame-perfect equilibrium, we mean supported by a subgame-perfect *trigger-strategy* profile. In Section 6 it is shown that in an important sense the trigger strategy restriction is without loss of generality

is the textbook *maximal deviation profit* when the collusive price is p^*, and

$$\pi^c(p^*, n, \delta) = \frac{1}{n} p^* D(p^*)/(1-\delta) \tag{11}$$

is the textbook *collusive profit* to each firm, at the collusive price p^*.

Moreover, if the revenue function $R(p) = pD(p)$ is single-peaked and p^* does not exceed the monopoly price $p^m \in \arg\max_{p\geq 0} pD(p)$, then a deviating firm wants to undercut the going price only slightly, that is, then $\hat{\pi}^t(p^*) = p^* D(p^*)$. Inequality (9) thus boils down to $1/n \geq 1-\delta$, so a collusive price p^* is then supported by a subgame perfect trigger-strategy equilibrium if and only if $\delta \geq 1 - 1/n$.

4. Analysis

There are two consequences of forward-looking intertemporal consumer behavior. First, a price-cut does not necessarily induce new consumers to buy since they might choose to wait for an even lower price in the future; this reduces the attractiveness of deviating from a collusive price. On the other hand, a price-cut can reel in old consumers who have not yet purchased the good at the high prices that prevailed in the past; this increases the attractiveness of deviation. The relative size of these two countervailing forces determines the payoff to deviation. We shall compare that payoff to the analogous deviation payoff in the textbook model. The parameter of greatest importance will turn out to be consumer patience. When consumers are patient the collusive force will be seen to prevail; consumers will prefer to wait for the price war, thereby lowering the deviation payoff to undercutting, and thereby discouraging deviations from the collusive price in the first place. On the other hand, when consumers are impatient, the fact that there are old customers who add their demand to that of the newly born ones will be seen to imply that deviation payoffs are higher than in the textbook model. Hence collusion will then be harder to sustain in the current model than in the textbook model.

One other interesting conclusion emerges from the analysis and that is the importance of "sales." Recall that in the textbook model if a firm were to undercut a collusive price, say p^*, then it would do so by undercutting ever so marginally. The reason for that is that any price below p^* would guarantee the whole market for the deviant, and – if the revenue function is single-peaked and p^* does not exceed the monopoly price – the whole market's revenue is highest at p^*. In the current model, by contrast, the most profitable under-cutting price sometimes turns out to be a small fraction of p^*. This is the case if p^* is the monopoly price. The reason is that a marginal price cut will fail to attract any old consumers.[14] Only a significant price cut can bring in old consumers, and since they represent a positive fraction of the potential buyers, a deviant firm, when the collusive price is at or near the monopoly price, would find it worthwhile to capture their demand.

4.1. Deviation profits and the equilibrium condition. Under perfect foresight, that is, when all consumers correctly anticipate the price war that will follow upon any price deviation, a trigger-strategy profile constitutes a subgame-perfect equilibrium if and only if

$$pD\left(p^{*}, p, 0\right) \leq \lambda \frac{p^{*} D(p^{*})}{n\left(1-\delta\right)} \tag{12}$$

for all $p < p^*$, where we recall from equation (6) that $D\left(p^{*}, p, 0\right)$ is the demand at the current price of p in a market where the price p^* has prevailed in the past, been expected to prevail in the past, and consumers currently anticipate the price to be zero in the future. The expression on the left-hand side is thus the current-period profit to a firm that unilaterally deviates to some price $p < p^*$. Such a deviation triggers a price war, so all future profits are then zero. The right-hand side is the present value of the stream of future profits to each firm, from the present period onwards, when there is no deviation. In such a situation, the price p^* is the market price in all periods – past, present and future – so then $p = p^e = p^*$. Hence, only new consumers buy in equilibrium.

For what follows it is useful to rewrite (12) by dividing both sides of the inequality by λ. That normalization allows us to write profits as profits *per new consumer* (recall the mass of new consumers is λ). By continuity of the left-hand side of (12) with respect to p, that inequality holds for all $p < p^*$ if and only if[15]

$$\max_{p \in [0, p^*]} \frac{p}{\lambda} D\left(p^{*}, p, 0\right) \leq \frac{p^{*} D(p^{*})}{n\left(1-\delta\right)}. \tag{13}$$

Let $\hat{\pi}\left(p^{*}, \beta, \lambda\right)$ denote the *maximal deviation profit* (per new consumer), the left-hand

[14] Recall that the only ones still in the market are those whose valuations are p^* or less. At a price less than but close to p^* these consumers are few and have virtually zero surplus from buying today but a positive surplus from buying next period (at a price of zero).

[15] The continuity of $D\left(p^{*}, p, 0\right)$ with respect to p follows from the continuity of $D = 1 - F$ and of the cut-off valuation $v^{+}\left(p, 0\right)$, which is well-defined and continuous in p for all $\beta < 1$.

side of (13). By equations (6) and (3), $\hat{\pi}(p^*, \beta, \lambda)$ equals

$$\max_{p\in[0,p^*]} p\left[D\left(\frac{p}{1-\beta}\right)+\frac{1-\lambda}{\lambda}\max\left\{0, D\left(\frac{p}{1-\beta}\right)-D\left(p^*\right)\right\}\right]. \tag{14}$$

When $\beta = 0$ and $\lambda = 1$, that is, when all consumers are maximally impatient and are in the market for only one period, then this maximal deviation profit is indeed identical with the maximal deviation profit in the textbook case.

We now turn to a study of the maximal deviation profit in order to determine what "collusive" prices p^* are sustainable in subgame-perfect trigger-strategy equilibrium. Of particular interest will be the comparison of the maximal deviation profit in this model with that in the textbook model. We note that the deviation profit function $\hat{\pi}$ is continuous, by Berge's Maximum Theorem and the assumed continuity of D. Moreover, it is non-increasing in consumer patience β. In other words, the more patient consumers are, the less does a deviating firm gain by under-cutting the others. We will therefore first consider the maximal deviation profit in the two extreme cases, when $\beta = 0$ and $\beta \to 1$. Then we will turn to intermediate β values.

4.2. Patient consumers. When consumers are maximally patient, $\beta \to 1$, new and old alike wait with gleeful anticipation for the price war that is to start in the period following a deviation, i.e., $D(p^*, p, 0) = 0$. The maximal deviation profit is accordingly zero. Hence, any price p^* for which demand $D(p^*)$ - and hence profit - is positive is sustainable as a trigger strategy equilibrium (see equation (13)). Note that p^* is sustainable as an equilibrium price no matter what the level of firm patience, δ, is. This is clearly in contrast with the textbook case.

By continuity, the deviation incentive remains negative for all β sufficiently close to 1. In other words, sufficiently patient consumers with perfect foresight are liable to face collusive prices. Since, moreover, the maximal deviation profit is monotonic in consumer patience, β, we thus have:

Proposition 1. *If $\delta \in (0, 1)$ and $p^* > 0$, then there exists a critical consumer discount factor, $\bar{\beta} < 1$, above which p^* is a subgame-perfect trigger-strategy equilibrium price and below which it is not.*[16]

How does this conclusion compare with the textbook model? Recall that there a collusive price p^* is subgame-perfect if and only if $\delta \geq 1 - 1/n$. According to

[16]Of course the critical consumer discount factor $\bar{\beta}$ in general depends on λ, δ, n and D.

the above proposition, by contrast, collusion is possible in the present model even if $\delta < 1 - 1/n$, granted consumers are sufficiently patient.

4.3. Impatient consumers. Let again the firms' discount factor δ be fixed at some positive level below 1, but now consider the polar case when consumers are maximally *impatient*, $\beta = 0$. Then all new consumers, along with all old with valuations below p^*, buy from any under-cutting firm. It follows immediately from equation (14) that the maximal deviation profit (per new consumer) then exceeds the collusive industry revenue $R(p^*) = p^* D(p^*)$:

$$\begin{aligned} \hat{\pi}(p^*, 0, \lambda) &= \max_{p \in [0,p^*]} p\left(D(p)/\lambda - (1-\lambda)\, D(p^*)/\lambda\right) \\ &\geq \max_{p \in [0,p^*]} pD(p) \geq R(p^*). \end{aligned} \tag{15}$$

Hence, if the revenue function is single-peaked and p^* does not exceed the monopoly price p^m, then $\hat{\pi}(p^*, 0, \lambda) \geq p^* D(p^*)$, i.e., collusion is "more diffcult" in the present model than in the textbook model. Indeed, this inequality is strict if $D(p)$ is strictly decreasing for prices below p^*.

It follows that if a collusive price p^* is sustainable in equilibrium in our model, then it is sustainable in equilibrium in the textbook model as well. By monotonicity and continuity of the deviation incentive with respect to consumer patience β, the deviation incentive exceeds that in the textbook model for all β sufficiently close to zero (see equation (9)). Consequently, collusion against impatient consumers is (weakly) harder to sustain in the present model than in the textbook model.

Proposition 2. *If $\delta, \lambda \in (0,1)$, $p^* \leq p^m$ and R is single-peaked, then there exists a critical consumer discount factor, $\widetilde{\beta} > 0$, below which p^* is a subgame-perfect trigger-strategy equilibrium price in the textbook model if it is a subgame-perfect trigger-strategy equilibrium price in the present model.*

4.4. Intermediate impatience. Having considered the two limiting cases when consumers are maximally and minimally patient, respectively, let us now consider intermediate cases, that is when $\beta, \delta, \lambda \in (0,1)$.

First note that, for any collusive price p^* not exceeding the monopoly price p^m, the maximal deviation profit, defined in equation (14), can be written as the largest of two deviation profits, one being the maximum deviation profit when the under-cutting price p does not exceed $(1-\beta)\, p^*$, and the other being when it does. In the first case, both new and old "bite", see equation (14), while in the second case only the new "bite." We will refer to the first case as a (unilateral) *sale* - a sizeable price cut in order to capture residual demand (c.f. Sobel (1984, 1991)). Formally,

$$\hat{\pi}(p^*, \beta, \lambda) = \max\{\hat{\pi}_1(p^*, \beta, \lambda), \hat{\pi}_2(p^*, \beta, \lambda)\}, \tag{16}$$

where

$$\hat{\pi}_1(p^*, \beta, \lambda) = \frac{1-\beta}{\lambda} \max_{p \in [0,(1-\beta)p^*]} \frac{p}{1-\beta}\left[D\left(\frac{p}{1-\beta}\right) - (1-\lambda)D(p^*)\right] \tag{17}$$

and

$$\hat{\pi}_2(p^*, \beta, \lambda) = (1-\beta) \max_{p \in [(1-\beta)p^*, p^*]} \frac{p}{1-\beta} D\left(\frac{p}{1-\beta}\right). \tag{18}$$

¿From now on, we focus on the case when the collusive price is the monopoly price. For $p^* = p^m$ we obtain that the second maximal deviation profit is $1-\beta$ times the maximal deviation profit in the textbook case, $\hat{\pi}^t(p^m) = R(p^m)$:[17]

$$\hat{\pi}_2(p^m, \beta, \lambda) = (1-\beta) R(p^m). \tag{19}$$

Hence, for any degree of consumer patience $\beta < 1$, such deviations are less profitable than in the textbook model.

What about more sizeable price cuts, that also make the old "bite"? For $p^* = p^m$ we evidently have

$$(1-\beta) R(p^m) \leq \hat{\pi}_1(p^m, \beta, \lambda) \leq \frac{1-\beta}{\lambda} R(p^m). \tag{20}$$

Hence,

$$\hat{\pi}_2(p^m, \beta, \lambda) \leq \hat{\pi}_1(p^m, \beta, \lambda) \qquad \forall \beta, \lambda \in (0,1), \tag{21}$$

that is, the maximal deviation profit can always be obtained in the lower price interval, where both the new and old "bite," irrespective of consumer patience and period length, as long as these lie strictly between zero and one. The reason that this is optimal, rather than to marginally undercut the going price, even when the population

[17] To see this, note that then $p \in [(1-\beta)p^*, p^*]$ is equivalent with

$$p^m \leq \frac{p}{1-\beta} \leq \frac{p^m}{1-\beta}$$

where the middle term is the maximand.

share λ of new consumers is close to 1 (its textbook value), is consumer patience: among the new consumers, only those with valuations above $p/(1-\beta)$ will buy now rather than wait for the ensuing price war.

In other words, optimal deviations always take the form of a "sale" rather than, as in the textbook model, a minute price cut. Moreover, if the demand function D is continuously differentiable, then such a profit-maximizing sales price p^s necessarily satisfies the first-order condition

$$R'\left(\frac{p}{1-\beta}\right) = (1-\lambda)\, D\left(p^m\right). \tag{22}$$

The left-hand side is continuous in p, ranging from a positive value exceeding the right-hand side when p is close to zero (its limit as $p \to 0$ is $D(0) > D(p^m)$) to zero as p approaches its upper bound $(1-\beta)\, p^m$.

Let p^s be any solution to equation (22). Then $p^s < (1-\beta)\, p^m$, and a necessary and sufficient condition for the monopoly price to be sustainable as the collusive price in subgame-perfect equilibrium is then

$$p^s \left(D\left(p^m\right) + \frac{1}{\lambda} \left[D\left(\frac{p^s}{1-\beta}\right) - D\left(p^m\right) \right] \right) \leq \frac{p^m D(p^m)}{n\,(1-\delta)}. \tag{23}$$

If we were to replace p^s by $(1-\beta)\, p^m$ in this inequality, then the quantity on the left-hand side would diminish to $(1-\beta)\, p^m D\left(p^m\right)$. Hence, a *necessary* (but in general not sufficient) condition for subgame perfection is

$$(1-\beta)\,(1-\delta) \leq \frac{1}{n}, \tag{24}$$

to be compared with the corresponding (necessary and sufficient) condition in the textbook case, namely $(1-\delta) \leq 1/n$. Likewise, the left-hand side is clearly less than $p^s D\left[p^s/(1-\beta)\right]/\lambda$, which in its turn is less than $(1-\beta)\, p^m D(p^m)/\lambda$. Hence, a *sufficient* condition for subgame perfection is

$$(1-\beta)\,(1-\delta) \leq \frac{\lambda}{n}. \tag{25}$$

We also note that if the revenue function is strictly concave, then the first-order condition defines a unique optimal sales price p^s in the open interval $(0, (1-\beta)\, p^m)$. This optimal sales price is a continuous and increasing function of consumer patience β and period length λ, and it tends to the monopoly price as $\beta \to 0$ and $\lambda \to 1$, so we do obtain the textbook deviation pricing rule as a limiting case.

In sum:

Proposition 3. *If $\beta, \delta, \lambda \in (0,1)$ and $p^* = p^m$, then the unilateral deviation profit is maximized at a price $p^s < (1-\beta)\, p^m$. If D is continuously differentiable, then such a deviation price satisfies equation (22), a necessary condition for which is (24) and a sufficient condition for which is (25). If R is strictly concave, then the optimal deviation price is unique, continuous, decreasing in β and increasing in λ, with limit value p^m as $(\beta, \lambda) \to (0,1)$.*

5. Example

Suppose valuations are uniformly distributed on the unit interval: $F(p) = p$ for all $p \in [0,1]$. Then $D(p) = \max\{0, 1-p\}$, $R(p) = \max\{0, p(1-p)\}$, and the monopoly price is $p^m = 1/2$. Consider collusive prices $p^* \leq p^m$, period lengths $\lambda > 0$, and consumer discount factors $\beta < 1$. A firm that undercuts by asking a price $p < p^*$ earns the deviation profit

$$\pi(p) = p\left(1 - \frac{p}{1-\beta} + \frac{1-\lambda}{\lambda} \max\left\{0, p^* - \frac{p}{1-\beta}\right\}\right). \tag{26}$$

Figure 1: Deviation profits as functions of the deviator's asking price.

Figure 1 shows the graphs of the function π when the collusive price is the monopoly price. The parabola with maximum at $p^m = 1/2$ represents the text-book case, and the humpy curve represents the case $\lambda = 0.1$ and $\beta = 0.75$. We see that the maximal deviation profit is smaller in the second case: instead of earning virtually $\pi = 0.25$ by marginally undercutting the collusive price, a deviating firm then can earn at most $\pi \approx 0.18$, and this is achieved by a substantial price cut, a "sale," in order to catch the residual demand from old consumers.

Figure 2 shows the deviation profit function for different degrees of consumer patience β, the middle curve representing the same case as in Figure 1, $\beta = 0.75$. The lower curve represents more patient consumers, $\beta = 1 - \lambda = 0.9$, and the higher curve less patient consumers, $\beta = 0.6$. In the last case the maximal deviation profit exceeds that in the textbook case.

Figure 2: Deviation profits as functions of the deviator's asking price, for different degrees of consumer patience β.

Figure 3 shows the deviation profit function for different consumer period lengths λ, other parameters being fixed, the middle curve representing the same case as in Figure 1, $\lambda = 0.1$. The taller curve represents a shorter market period (and hence larger share of old consumers), $\lambda = 0.05$, and the lower curve represents a longer market period, $\lambda = 0.25$. With the shorter market period, the maximal deviation profit again exceeds that in the textbook case.

Figure 4 shows how the function π changes as we approach the textbook model, beginning with the same parameters β and λ as in Figure 1: the parameter sequence being $(\beta, \lambda) = (0.75, 0.1)$, $(0.5, 0.5)$ and $(0.25, 0.75)$. The graph of the deviation profit function point-wise approaches the graph of the textbook deviation profit function $(\pi(p) = p(1-p))$, while the optimal deviation price remains, as was shown analytically in the preceding section, in the low price interval where both new and old consumers bite.

Figure 3: Deviation profits as functions of the deviator's asking price, for different period lengths λ.

Figure 4: Deviation profits as functions of the deviator's asking price, for different combinations (β, λ) of consumer patience and period length.

We finally consider the optimal deviation price p^s. In the present example, the defining equation (22) boils down to

$$p^s = \frac{1}{4}(1-\beta)(1+\lambda). \tag{27}$$

Figure 5 shows isoquants for the optimal sales price p^s, as a function of the period length λ and consumer patience β. As these two parameter values move away from their textbook values, the optimal deviation profit falls. Of particular interest is the case when consumers' discount factor β is coupled to the period length λ. One such case is when this discount factor equals the "survival probability" $1-\lambda$ from one market period to the next. This corresponds to the negative-sloped straight line in Figure 5. A more general case is when the discount factor β is the product of the survival probability $1-\lambda$ and an exponential pure time-preference discount factor, such as $\beta = (1-\lambda)\exp(-r\lambda)$ for some (subjective) discount rate $r \geq 0$. This corresponds to the negatively sloped curve in Figure 5 (drawn for $r = 1$).

Figure 5: Isoquants for the optimal sales price, with the market-period length λ on the horizontal axis and consumer patience β on the vertical.

In Figure 6, finally, is plotted the curve in the (λ, β)-plane, where the optimal deviation profit in the present model,

$$\pi^s = \frac{(1-\beta)(1+\lambda)^2}{16\lambda}, \tag{28}$$

equals $1/4$, the supremum deviation profit in the textbook model. The area above the curve are parameter combinations (λ, β) at which collusion is easier in the present model than in the textbook model ($\pi^s < 1/4$), while for parameter combinations below the curve collusion is harder than in the textbook case. In particular, if consumers' discount factor β equals their survival probability from one market period to the next, that is, if $\beta = 1-\lambda$, then collusion is easier in the present model than in the textbook model for all $\lambda \in (0, 1)$.

Figure 6: The 1/4-isoquant for the optimal sales profit.

6. Equilibria in Other Strategies

So far we have focused exclusively on the use of "grim" 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 grim trigger strategies.

Consider first "forgiving" trigger strategies with a finite punishment horizon, that is, a punishment strategy in which firms price at zero during the next T periods after a deviation and thereafter revert to the collusive price p^*, for some positive integer T.[18] Seeing a deviation price $p < p^*$ when firms use such strategies, consumers have two relevant options: either buy immediately at price p, wait a period and buy it at price zero, or plan not to buy at all. Evidently it does not pay to plan to wait any longer since subsequent prices are nonnegative. The consumer decision problem is thus identical to the one under grim trigger strategies, and hence the aggregate demand function is identical in the two cases. In turn that implies that the maximal deviation profit in the current period is identical. Future profits after a deviation are, however, higher since the deviation is "forgiven" after T periods. In sum: (total) deviation profits are higher under forgiving trigger strategies than under grim trigger strategies. Hence, the set of sustainable collusive prices is a subset of that under

[18] In principle one needs to consider deviations both from the collusive price and from punishment of a deviator. However, in the present case, no profitable deviation exists in a punishment phase.

grim trigger strategies.

Consider now any general punishment strategy supporting a collusive price p^*. The implied punishment is necessarily milder than under a grim trigger strategy. A consumer with perfect foresight now anticipates the path of prices $p^e(\tau)$, $\tau = 1, 2, ...$ 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 $v \geq p$ and

$$v - p \geq \max_{\tau} \beta^{\tau} \left[v - p^e(\tau) \right]. \tag{29}$$

It immediately follows that the quantity on the right-hand side is not larger than under the grim trigger punishment.[19] Hence, the incentives of consumers to wait is the highest under the grim trigger strategy. Put differently: the current period deviation profit is (weakly) larger for a deviating firm under any other punishment strategy. This is a new phenomenon that does not exist in the textbook model: intertemporal consumers are more likely to buy in the current period if they anticipate a milder price war. A second effect – which does exist in the textbook model – is that a milder punishment also implies that future profits are higher, as compared with those under the grim trigger strategy. Therefore, the two effects work in the same direction, and together imply that (total) deviation profits are higher under any alternative punishment strategy. Hence:

Proposition 4. *Any price p^* that can be sustained in subgame-perfect equilibrium can be sustained in subgame-perfect equilibrium under the grim trigger strategy.*

7. Conclusion

Our model of Bertrand competition in recurrent market interaction is built on a number of simplifying assumption. 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. Such a generalization does not appear to be analytically feasible, in particular if β and v are statistically independent at the individual level. As an aside, we also note that all our results hold also if consumers have non-exponential time preferences, as long as these are strictly monotonic. We effectively only use the discount factor β from the current period to the next.

[19] Recall that under the grim trigger strategy $p^e(\tau) = 0$ for all $\tau > 0$.

Another simplification is that we have focused on the case of an indivisible durable good. It seems likely that the qualitative results carry over also to the case of divisible and non-durable goods – another task 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. 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. More generally, a study of collusion against boundedly rational consumers would be a valuable extension.

Another potentially interesting extension of the present analysis is to investigate the possibility of temporary sales on the equilibrium path. For while we only consider constant collusive prices, a temporary equilibrium sale now and then could be a way for the colluding firms to capture the residual demand in the market, much along the lines in Sobel (1984, 1991). While no consumer "dies" in Sobel's models, and therefore residual demand under a constant market price then builds up without bounds over time, our consumers do "die," and therefore residual demand is bounded. However, the residual demand may nevertheless be sufficiently large to motivate collusive temporary sales. One such pricing scheme, which seems potentially amenable to analysis, is when firms in equilibrium run a common "sale" every T periods, at some price $p^{**} < p^*$.[20] During such a sale, consumers with perfect foresight will anticipate the market price to revert to its "normal" higher level next period. Hence, even if they are patient, they cannot gain by waiting. Thus, aggregate demand in such an equilibrium sale can be much higher than under a unilateral price deviation, as studied above. On the other hand, the cut-off valuation in "normal" periods will be lower than under a constant collusive price, since consumers anticipate the next equilibrium sale. Thus, firms' profits in "normal" periods will be somewhat lower than under a constant collusive price. The net effect of these two opposing forces will determine whether or not equilibrium sales are attractive collusive devices. We here refer to Sobel's work, and leave this for future research.

We finally mention yet another avenue for future research, namely to apply the

[20]If the market periods are months, then the firms may, for example, agree to have a common sale every January.

present model of intertemporal demand in analyses of dynamic Cournot oligopoly.

References

[1] Ausubel L. and R. Deneckere (1987): "One is almost enough for monopoly", *Rand Journal of Economics* 18, 255-274.

[2] Coase R.H. (1972): "Durability and monopoly", *Journal of Law and Economics* 15, 143-149.

[3] Dutta, P. K. (1995): "A Folk theorem for stochastic games", *Journal of Economic Theory* 66, 1-32.

[4] Fudenberg, D. and J. Tirole (1991): *Game Theory*, MIT Press, Cambridge, MA.

[5] Gul F., H. Sonnenschein and R. Wilson (1986): "Foundations of dynamic monopoly and the Coase conjecture", *Journal of Economic Theory* 39, 155-190.

[6] Gul F. (1987): "Noncooperative collusion in durable goods oligopoly", *Rand Journal of Economics 18*, 248-254.

[7] Maskin, E. and J. Tirole (1988): "A theory of dynamic oligopoly, II: Price competition, kinked demand curves and fixed costs", *Econometrica* 56, 571-600.

[8] Kirman A. and M. Sobel (1974): "Dynamic oligopoly with inventories", *Econometrica* 42, 279-287.

[9] Radner, R. (1999): "Viscous demand", mimeo., Stern School, NYU, New York.

[10] Selten R. (1965): "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit", *Zeitschrift für die gesamte Staatswissenschaft* 121, 301-324, 667-689.

[11] Sobel J. (1984): "The Timing of Sales", *Review of Economic Studies* 51, 353-368.

[12] Sobel J. (1991): "Durable Goods Monopoly with Entry of New Consumers", *Econometrica* 59, 1455-1485.

[13] Tirole J. (1988): *The Theory of Industrial Organization*, MIT Press, Cambridge, MA.

[14] Wallner, K. (1999): "Sequential moves and tacit collusion: Reaction function cycles in a finite pricing duopoly", *Journal of Economic Theory*, 84, 251-267