Product Differentiation, Uncertainty and the Stability of Collusion

Michael Raith*

ECARE

Université Libre de Bruxelles

July 29, 1996

---

* Correspondence: ECARE, Université Libre de Bruxelles, 39 Av. Franklin Roosevelt, 1050 Bruxelles, Belgium; email: mraith @ ulb.ac.be. This paper originates from Chapter 4 of my Ph.D. thesis at the London School of Economics (Raith 1996). I would like to thank John Sutton for his guidance throughout the paper's evolution, and Patrick Bolton and especially Leonardo Felli and Paul Klemperer for their very constructive suggestions. For comments on this or previous versions, I am also indebted to Guido Friebel, Godfrey Keller, Sougata Poddar, George Symeonidis, and seminar participants at various occasions. Any remaining errors are my own. Financial support by the Suntory-Toyota International Centres for Economics and Related Disciplines (at LSE) and by the European Commission (through its Human Capital and Mobility Programme) is gratefully acknowledged.

Abstract

The conventional view that product heterogeneity limits the scope for collusion among oligopolists has been challenged in recent theoretical work. This paper provides an argument in support of the conventional view by emphasising the role of uncertainty. I introduce the idea that, with stochastic demand, an increase in the heterogeneity of products leads to a decrease in the correlation of the firms' demand shocks. With imperfect monitoring, this makes collusion more difficult to sustain, as discriminating between random demand shocks and *marginal* deviations from the cartel strategy becomes more difficult. These effects are illustrated within a Hotelling-type duopoly model.

JEL-Codes: D43, D82, L41

Keywords: Collusion, Product Differentiation, Imperfect Monitoring

1 Introduction

Does heterogeneity of products limit the scope for collusion among oligopolists, or is it rather a facilitating factor? Many economists would hold the former to be true, arguing that heterogeneity, in some sense, makes the firms' coordination problem more complex. A recent surge of interest in this question[1], however, has led to numerous game-theoretic models which, owing to their particular specifications, predict that cartel stability *increases* with the degree of product differentiation.

In contrast, this paper argues that uncertainty, neglected in both the informal literature and the new theoretical contributions, alters the problem of sustaining collusion in a fundamental way and plays in a crucial role in determining the effect of product differentiation on the scope for collusion. I introduce the idea that an increase in the heterogeneity of products leads to a decrease in the correlation of the demand functions for the goods. In an environment where a firm cannot observe its rivals' actions but has to infer from observable signals whether another firm has deviated, this in turn makes collusion more difficult to sustain, as discriminating between random demand shocks and *marginal* deviations from the cartel strategy becomes more difficult. These effects are illustrated in a modified version of the Hotelling-type model with taste heterogeneity due to de Palma et al. (1985). The results stand in sharp contrast to those in the recent theoretical literature. At the same time, the model provides a simple analytical foundation for the traditional view that heterogeneity limits the scope for collusion.

According to the traditional view, heterogeneity impedes cartel behaviour because, in some sense, firms face a situation of "higher complexity". For example, while with homogeneous products firms merely have to agree on one price, with heterogeneous

[1] For horizontal product differentiation focused on here, this includes the papers by Deneckere (1983), Wernerfelt (1989), Chang (1991), Ross (1992) and Häckner (1993), and the discussion in Martin (1993, p. 116). The case of vertical product differentiation is studied by Häckner (1994). In contrast, the papers by Jehiel (1992) and Friedman and Thisse (1993) study the effect of collusive conduct on firms' effort to differentiate products, and Zhang (1995) can also be counted to this strand of literature.

goods a whole array of prices has to be negotiated. This problem "grows in complexity by leaps and bounds" (Scherer and Ross 1990, p. 279) with the number of characteristics in which the goods can differ. Similarly, firms may have difficulty in monitoring the policies of their rivals in complex situations (Clarke 1985, p. 60).[2]

Though intuitively compelling, it is difficult to pin down analytically an appropriate interpretation of this argument. For example, one could argue that with heterogeneous goods, the relevant space of product attributes becomes very large and may not even be specifiable in advance, and thus would render both cartel negotiations and subsequent enforcement increasingly difficult. This may be true, but will be difficult to capture in standard models of product differentiation, since in these models products are usually symmetrically positioned in a relatively simple space of characteristics. Here, it is harder to see in which sense differentiation could lead to a situation of increased complexity.[3] Put differently, if the intuition that heterogeneity has something to do with complexity is correct, then product differentiation has implications not captured by the standard models.

The traditional view that heterogeneity hinders collusion has been recently challenged. Following Deneckere (1983), several theorists have analysed the effect of product differentiation on cartel stability within game-theoretic models. All models that have been studied are deterministic models in which deviations from a cartel strategy are detected immediately and precipitate retaliation. Therefore, sustainability of collusion only depends on the tradeoff between the benefit from collusion (equivalently, the

[2] Posner (1976, p.60) is more precise, but also makes specific explicit assumptions about the information structure and the nature of cartel agreements: "... the detection of cheating by members of the cartel will be complicated by the difficulty of knowing whether a competitor's price is below the agreed level or is simply a lower price for a lower grade or quality of the product."

[3] Similarly, Tirole (1988, p.240) notes that while the role of detection lags as a factor hindering collusion is well understood, "efforts to formulate the second factor [asymmetries, as which he counts the case of differentiated products] have not been as successful".

severity of punishment), and the gain obtained by deviating from a cartel strategy. In this literature, cartel stability is measured by the critical discount rate below which the joint profit maximising price can be sustained with a trigger strategy (in most cases a simple grim trigger strategy; for an exception, see Wernerfelt [1989]). For this case, the critical discount rate is simply the ratio of the collusive benefit to the defection gain (cf. Martin 1993, p.104). A decrease in the critical discount rate is then interpreted as a decrease in the scope for collusion.

Under these circumstances, both the benefit of collusion and the gain of deviating from a cartel strategy are likely to be smaller for differentiated goods than for homogeneous goods. Hence, there would be little reason to expect any systematic relationship between the critical discount rate and the degree of product differentiation, and indeed there are several results showing an ambiguous relationship.[4] Nevertheless, there are some models in which the critical discount rate *increases* as products become more differentiated. This has led some authors to question the validity of the traditional view that heterogeneity hinders collusion (cf. Ross 1992), and has even led to the emergence of a new conventional wisdom among theorists.

A serious problem with this research programme is the disregard of any uncertainty which might play a role. In particular, it is implicitly assumed that any deviation from collusive behaviour is detected with certainty. With this information structure, firms can only choose between two extremes: either to adhere to the cartel strategy, or to cut the price (or increase the quantity) by a large amount so as to maximise the current-period profit, in anticipation that retaliation will follow with certainty. In contrast, as Stigler (1964) argued, one would more realistically expect firms to consider increasing their profits by cutting their price only *slightly*, in the hope that such a deviation will go unnoticed by the other firms. In addition, even if firms have excess capacity, capacity constraints will, in the short run, in general not allow firms to take over substantial market shares from other firms in the first place.

Empirical evidence on the relationship between product differentiation and cartel

[4] E.g. Deneckere (1983), Ross (1992), Wernerfelt (1989).

stability is usually drawn from antitrust sources.[5] Here, the conventional wisdom that homogeneity eases collusion may induce a "selection effect": the authorities may focus on homogeneous goods industries.[6]

A different and perhaps more convincing kind of support for the traditional view can be found in case studies which emphasise the role of strategic product *standardisation*. For example, in his study of the US electrical industry, Sultan (1974, p.28-29) points out that the main purpose of an "organized industry effort to standardize the designs of most products" during the 1920s was to reduce price warfare by increasing the visibility of price-cutting.

Overall, the available evidence for the traditional view that heterogeneity hinders collusion is rather weak. On the other hand, I am not aware of any evidence supporting the opposite conjecture that differentiation facilitates collusion.[7]

In this paper, I add a new dimension to the analysis of cartel stability, viz. uncertainty and, hence, the *probability* of price wars being triggered in the first place. The static model is a variant of the Hotelling-type model due to de Palma et al. (1985). In Section 2, it is shown how a combination of random "macro" shocks on the demand side, affecting the density of consumers along the line, and heterogeneity in tastes among consumers, generate positively correlated demand functions for the two goods, such that the correlation coefficient depends positively on the degree of substitutability between the goods. Intuitively, if the two goods are similar, they attract the same groups of customers. Shocks on the demand side, are then reflected in the firms' demand functions in largely the same way. On the other hand, if goods are differentiated,

[5] cf. Hay and Kelley (1974), Fraas and Greer (1977).

[6] In fact, Hay and Kelley (1974) find that "virtually all of the entries [in their sample] would read 'high' product homogeneity."

[7] This assertion is consistent with an important observation made by Fraas and Greer (1977): where the environment is most conducive to collusion, firms can collude tacitly, whereas under less favourable conditions firms may require formal cartel arrangements to sustain collusion. In this sense, collusion will tend to be most visible when colluding is more difficult.

the lower is the correlation of the demand functions because the firms attract customers from different populations.

In Section 3, we analyse collusion between the two firms. Each firm cannot observe either the shocks or the other firm's price, but can observe its own and the rival's realised demand, and infers from these whether the other firm has deviated from the cartel strategy. Firms collude using trigger strategies of a generalised Green-Porter (1984) type. With the uniform distribution of the density shocks used here, optimal strategies in this class take a particularly simple form: a price war is only triggered if one firm has deviated for sure, i.e. when a quantity vector is observed which could not be observed if both firms adhered to the cartel price. In this case, firms play the static Bertrand equilibrium forever. On the other hand, a deviation is of course not necessarily detected.

We determine the minimum discount factor above which collusion can be sustained using a trigger strategy. It turns out that this critical discount factor is decreasing in the degree of substitutability of the goods. In this sense, collusion is *less* sustainable the more differentiated the goods are.

The effect of product differentiation on the correlation of the demand functions decreases the scope for collusion, since with a lower correlation, discriminating between cheating and exogenous random fluctuations in demand is more difficult. More precisely, it is the *increase* of the probability of retaliation if a firm deviates that matters for the sustainability of a certain strategy. Now, the more differentiated the products are, the lower is the correlation, and therefore the smaller is the effect of a deviation on the probability of a price war. So the retaliation phase loses its deterrent effect, which undermines the stability of the cartel.[8]

[8] More precisely, while the set of events that trigger a price war is itself endogeneous, adjusting this set cannot cancel the described effect. Rather, cartel stability is undermined because with more heterogeneity, optimally designed strategies must be more lenient than in the case of homogeneous goods.

2 The static model

In this section, I discuss the stage game of the model. In particular, I will show how demand shocks give rise to demand functions, the correlation between which depends on the degree of product differentiation.

The model is a modified version of the well-known Hotelling model with taste heterogeneity due to de Palma et al. (1985). Two firms are located symmetrically on a line $[0,1]$. Firm 1 is located at $\sigma/2$ and firm 2 at $1-\sigma/2$, where $0 \leq \sigma \leq 1$. Thus σ measures the degree of product substitutability: for $\sigma = 1$, we have homogeneous goods, and for $\sigma = 0$, maximal differentiation.

Consumers are distributed along the line. The utility of a consumer located at z who purchases firm i's product is given by

$$w_i(z) = y + a - p_i - (z - z_i)^2 + \varepsilon_{iz},$$

where y denotes income, a is the utility derived from consuming the most preferred good, p_i is firm i's price and z_i its location (for simplicity, the parameter for the "travel costs" is set to 1).

Heterogeneity of tastes is introduced by means of the random variables ε_{1z} and ε_{2z}, which are assumed to be i.i.d. double exponentially distributed with zero mean and variance $\mu^2\pi^2/6$ (cf. Anderson, de Palma and Thisse 1992, p.363). In effect, through this random utility specification a second dimension of product differentiation (measured by the paramater μ) is introduced into the model, in addition to differentiation along the Hotelling line.[9] The significance of this specification, however, is that it leads to an overlapping of the firms' market areas in a natural way, as I will discuss in detail below.

With the distribution of ε_{iz} as specified above, the probability that consumer z buys at firm 1 is then given by the logistic function

$$P_1(z) = \left\{1 + \exp\left[\frac{1}{\mu}\left(p_1 - p_2 + (1-\sigma)(2z-1)\right)\right]\right\}^{-1}. \qquad (1)$$

[9] For an extensive discussion of this formalisation of product differentiation, see Anderson, de Palma and Thisse (1992).

In contrast to de Palma et al. (1985) and other models, I assume there are "macro" shocks affecting the density of consumers in different parts of the Hotelling line. More specifically, the consumers fall into two groups, group 1 in the interval $[0, 1/2)$ and group 2 in $(1/2, 1]$. The densities of consumers in both intervals, u_1 and u_2, are independent and uniformly distributed over $[1-d, 1+d]$, with $d \in [0, 1]$.[10] These density shocks can be thought of as being caused by taste changes, business cycles, or other reasons, affecting groups 1 and 2 in different ways. Other things equal, consumers in group 1 have a preference for good 1, and similarly for group 2. But with taste heterogeneity, there is always a positive probability that a consumer will purchase the "other" product, where this probability depends both on prices and the similarity of the products. This is obviously a very crude and simple way of introducing uncertainty into the model; many other specifications are conceivable. The parameter d measures the degree of demand uncertainty. Maximal uncertainty obtains if $d = 1$, whereas for $d = 0$, the model converges to the deterministic model studied by de Palma et al. (1985).

The description of demand has two salient features. First, the random utility approach used here implies a "cross-over" of market areas[11]; i.e., even with differentiation, both firms attract consumers from the entire Hotelling line. Second, the density shocks (i.e. market size shocks) in both halves of the line imply that the correlation of the firms' demands depends on the proportions of customers each firm draws from each half of the line, and therefore depends on the degree of product differentiation.

The firms' demand functions q_1 and q_2 are obtained by integrating the purchase probabilities $P_1(z)$ and $1 - P_1(z)$, respectively, over the entire line, taking the densities u_1 and u_2 into account:

$$q_1(k, \mathbf{u}) = q_1^l(k)u_1 + q_1^r(k)u_2 \qquad \text{and} \qquad q_2(k, \mathbf{u}) = q_2^l(k)u_1 + q_2^r(k)u_2, \tag{2}$$

[10] The assumption of independence simplifies calculations but is not necessary; any degree of positive correlation would leave the qualitative features of the results unaffected.

[11] Cf. Archibald et al. (1986) and de Palma et al. (1994).

where $k = p_1 - p_2$, $\mathbf{u} = (u_1, u_2)$ and

$$\begin{aligned}
q_1^l &= \int_0^{1/2} P_1(z)\mathrm{d}z &=& \frac{\mu}{2(1-\sigma)} \log \frac{1+e^{(-k+1-\sigma)/\mu}}{1+e^{-k/\mu}}, \\
q_1^r &= \int_{1/2}^{1} P_1(z)\mathrm{d}z &=& \frac{\mu}{2(1-\sigma)} \log \frac{1+e^{-k/\mu}}{1+e^{(-k-1+\sigma)/\mu}}, \\
q_2^l &= \int_0^{1/2} P_1(z)\mathrm{d}z &=& \frac{\mu}{2(1-\sigma)} \log \frac{1+e^{k/\mu}}{1+e^{(k-1+\sigma)/\mu}}, \\
q_2^r &= \int_{1/2}^{1} P_1(z)\mathrm{d}z &=& \frac{\mu}{2(1-\sigma)} \log \frac{1+e^{(k+1-\sigma)/\mu}}{1+e^{k/\mu}}
\end{aligned}$$

(the superscrips 'l' and 'r' stand for 'left' and 'right'). The firms' demand functions are correlated random variables, since they are functions of the same random vector $\mathbf{u}$. Before we go on to analyse noncooperative and collusive pricing in this game, let us first look at what determines the correlation of the firms' demand levels. For simplicity, consider the case where both firms charge equal prices, i.e. $k = 0$. From (2) we have $q_1 = u_1 m_1 + u_2 m_2$ and $q_2 = u_1 m_2 + u_2 m_1$, where

$$m_1 = \frac{\mu}{2(1-\sigma)} \log \frac{1+e^{(1-\sigma)/\mu}}{2} \quad \text{and} \quad m_2 = \frac{\mu}{2(1-\sigma)} \log \frac{2}{1+e^{-(1-\sigma)/\mu}}.$$

Since u_1 and u_2 are independent, the correlation between q_1 and q_2 is then given by $\rho_u = 2m_1 m_2/(m_1^2 + m_2^2)$. How this correlation depends on the substitutability parameter σ and on μ is depicted in Figure 1.

For σ close to one, the correlation is close to unity: since the products are virtually identical, consumers buy either of the two goods with almost equal probability. On the aggregate level, the demand functions become more or less identical (i.e., both m_1 and m_2 are close to 1/4). In contrast, for larger degrees of differentiation, group-1 consumers will buy good 1 with a higher probability than good 2, and vice versa for group-2 consumers. On the aggregate level, this implies that firm 1 will draw its customers to a larger extent from group 1 than from group 2 (formally, $m_1 > m_2$), and therefore will be more affected by the shock m_1 than by m_2. Since the reverse holds for firm 2, this implies a lower correlation between the two firms' demand functions. Thus, the model precisely captures the idea discussed in the Introduction that the correlation of demand functions depends on the degree of product differentiation.

Figure 1: Correlation ρ_u between q_1 and q_2.

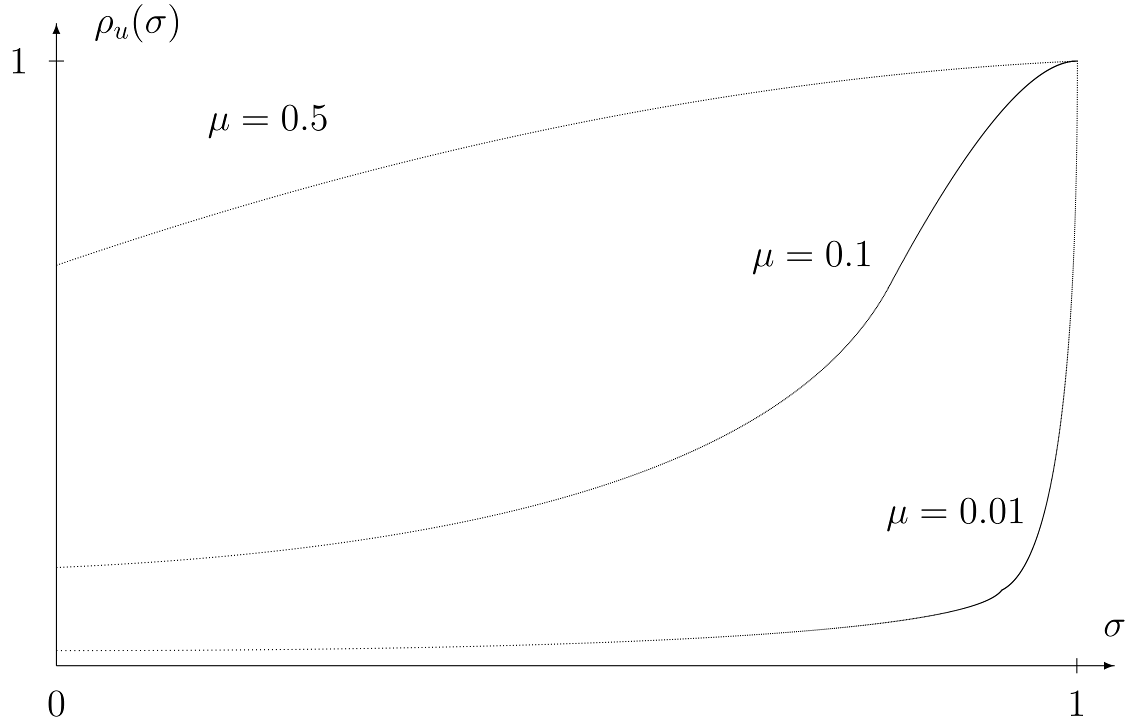

The correlation is increasing in the heterogeneity parameter μ: for small μ, i.e. little heterogeneity, even slight differentiation along the Hotelling line will suffice to effectively separate the firms' market areas into the two halves of the line. In contrast, with large μ, the correlation of the firms' demand functions will be considerable even with maximal differentiation along the Hotelling line.

The degree of product differentiation affects not only the correlation of the firms' demand functions, but also their variance: with heterogeneous goods, the variance of each firm's demand is larger than with more similar goods. This relationship, too, seems economically very plausible: a firm that targets a specific group of customers is mainly affected by variations of demand by that particular group, whereas a firm that offers a standard product purchased by different customers is less vulnerable to shocks affecting a particular group.

I follow Green and Porter (1984) and related papers in the tradition of Stigler (1964) in assuming that in each period, firms set prices before the shocks u_1 and u_2 are realised. Hence, the firms' pricing decisions are determined by their expected demand, which is obtained by setting $u_1 = u_2 = 1$ in (2).

Assume that the firms produce at constant marginal cost, which in this Hotelling framework can be set to zero without loss of generality. Then firm 1's expected profit is given by $\pi_1 = p_1[q_1^l(k)+q_1^r(k)]$. Firm 2's expected profit π_2 is defined analogously. Given the symmetric locations of the firms, it is easy to establish that the Nash equilibrium in prices is unique and symmetric. This Bertrand equilibrium price p^b is obtained by expressing π_1 as $(k + p^b)[q_1^l(k) + q_1^r(k)]$ and solving the first-order condition

$$\frac{\partial \pi_1}{\partial p_1} = \frac{\partial \pi_1}{\partial k} = q_1(k) + (k + p^b)\frac{\partial}{\partial k}[q_1^l(k) + q_1^r(k)] = 0$$

for p^b at $k = 0$, which gives

$$p^b = (1 - \sigma)\frac{e^{(1-\sigma)/\mu} + 1}{e^{(1-\sigma)/\mu} - 1}.$$

For $\mu \to 0$, this price converges to $(1 - \sigma)$, the equilibrium price in the standard model of d'Aspremont et al. (1979). Moreover, p^b increases with μ, reflecting a larger degree of product differentiation due to more taste heterogeneity.

3 Collusion

I now turn to the analysis of collusion. As argued by Stigler (1964), in a world of imperfect information, collusion among oligopolists is threatened by the possibility of secret price-cutting on part of some cartel members. A firm considering a deviation from an agreed cartel price faces a trade-off between additional profits to be gained, and the risk of detection and punishment of such price-cutting by the other cartel members. It is this trade-off, which relates to the effects of *marginal* price cuts, which I am going to analyse in this paper. The analysis is therefore closely related to papers which have formalised Stigler's argument, notably, Green-Porter (1984).

In contrast, in the literature initiated by Deneckere (1983) which analyses how cartel stability depends on product differentiation (cf. Introduction), there is no uncertainty, and therefore deviations from the cartel strategy are immediately detected and punished. Hence, firms deviate, if at all, not marginally, but by a large price cut (or quantity increase) in order to maximise their current-period payoff.

The analysis below shows that the sustainability of collusion with respect to marginal deviations leads to a restriction on the cartel price which, to some extent, also rules out the profitability of large price cuts. For rather homogeneous goods, however, it is not necessarily true that if marginal deviations are not profitable, large deviations are not profitable either: marginal deviations may be deterred by large marginal increases in the probability of detection. Still, it may be worthwhile to cut the price by a large amount and maximise the current-period payoff even if this leads to retaliation for sure. In this situation, for a collusive strategy to be sustainable, two incentive constraints must be satisfied, one for marginal and one for large price cuts.

Large price cuts are discussed in detail in Appendix A. Here, the procedure of analysis, and to some extent the predictions, are quite similar to those of other papers which deal with deterministic models. Hence, by simultaneously considering marginal and large price cuts, the analysis synthesises two strands of literature: deterministic models such as Deneckere (1983) and related works (cf. footnote 1), and models with imperfect information such Green-Porter (1984). I analyse the conditions under which the sustainability of collusion is constrained by marginal deviations rather than large price cuts. It turns out that there are three important determining factors: first, the level of uncertainty, parameterised by d, obviously matters for the relevance of marginal deviations, since for $d \to 0$, the model converges to a deterministic Hotelling model in which any deviation is detected with certainty. Second, the closer the collusive price is to the Bertrand price, the less profitable are large deviations, and hence the more relevant are marginal deviations. Finally, capacity constraints, too, limit the profitability of large price cuts.

3.1 Information structure

Following the literature in the tradition of Stigler (1964), I assume that a firm cannot observe its rival's price. Moreover, I assume that it can observe both its own and the rival's realised demand in each period.

Suppose, for a moment, that a firm could observe only its own demand but not the rival's (as Stigler assumed). Price wars might then be initiated by either of the firms on grounds of - not verifiable - unusually low sales due to cheating on part of the other firm. But given that own demand is private information, a firm would *ex post*, after observing a very low demand, never have an incentive to lead the firms into a costly price war, even if it were certain that the other firm had indeed deviated.[12] As a result, collusion could never be sustained in an ordinary Nash equilibrium. An analysis of such a game with imperfect private information would therefore require the use of a different equilibrium notion, and would lead to severe technical complications.[13] Such problems are circumvented if we assume that the occurence of price wars is conditioned on realisations of a *public* signal, in our case, the vector of quantities.

Moreover, an economic argument in favour of this assumption is that both for tacitly colluding firms and for organised cartels, shipped quantities are likely to be better observable than the accompanying monetary flows (according to Ulen's [1978, p.128] description of the Joint Executive Committee, this cartel made an effort to monitor both shipments and billing, although apparently cheating remained a possibility).

3.2 Optimal collusive strategies

The trigger strategies considered here have the following structure: in the first period, the game is in a "collusive" mode, in which both firms set a cartel price p^c. There is a set of quantity vectors T, called the trigger set. If in any period the realised quantity vector $\mathbf{q}$, which is observed by both firms, belongs to the trigger set, then a "punishment phase" is triggered, otherwise the game remains in the collusive mode. I do not make any assumptions on how exactly firms behave during such a punishment phase, apart from the basic requirement that the firms follow the path of some sequential equilibrium. All that is needed for our purposes is the assumption that the continua-

[12] Cf. the discussion in Fudenberg and Levine (1991).

[13] On the analysis of games with imperfect private information, see Fudenberg and Levine (1991) and Lehrer (1992). An oligopoly model with private but perfect information (arising due to localised competition) is studied by Verboven (1995).

tion payoff at the beginning of a punishment phase be smaller than the continuation payoff in the collusive mode. Formally, if the continuation payoff in collusive mode is v^c, let the continuation payoff at the beginning of a punishment phase triggered by the occurence of $\mathbf{q}$ be given by $\eta(\mathbf{q})v^c$, with $\eta(\mathbf{q}) < 1$ for all $\mathbf{q} \in T$. For simplicity I assume that the worst punishment is given by firms playing the Bertrand equilibrium indefinitely.

Thus, for the strategies considered here there is no claim of global optimality, in contrast to Abreu's (1986) stick-and-carrot strategies for deterministic models and the strategies analysed by Abreu, Pearce and Stacchetti (1986) for models with imperfect information. On the other hand, the strategy set here clearly includes the trigger strategies of Green and Porter (1984), which specify the same kind of punishment for each signal that triggers a retaliation phase.

It now turns out that if firms use strategies which maximise their expected discounted profit, then, due to the specific nature of the demand shocks introduced in the previous section, the optimal collusive strategies have a particularly simple structure. To state this result, denote by Σ the set of all possible quantity vectors (for any price vectors), and by $S(0)$ the support of $\mathbf{q}(\mathbf{u})$ in case both firms adhere to the cartel price p^c, i.e. where $k = 0$.

Proposition 1 *Any optimal collusive strategy is characterised by some cartel price p^c and the trigger set $T = \Sigma - S(0)$. Upon observing any quantity vector $\mathbf{q} \in T$, firms play the static Bertrand equilibrium forever.*

Proof: see Appendix B.

This result has two sides to it: first, it is clear that it is optimal to punish the occurence of *every* vector outside $S(0)$, i.e. one that can only be observed if a firm has deviated for sure. Moreover, since in equilibrium such events are never observed, it is optimal to punish them maximally. Thereby, maximal compliance is achieved without involving any cost for the cartel.

Second, according to Proposition 1, the occurence of a vector $\mathbf{q} \in S(0)$ should never trigger a punishment phase. This feature, which is in contrast to strategies considered by Green and Porter (1984) or Abreu, Pearce and Stacchetti (1986), results from the particular distribution of the demand shocks considered here: including any vectors in $S(0)$ in the trigger set is costly since this implies that a punishment phase occurs with positive probability. Therefore, this can be optimal only inasmuch as it serves to deter deviations from the cartel price. Any deterrence effect can only result from an *increase* in the probability of a price war in case a firm deviates. The proof of Proposition 1 now shows that due to the rectangular shape of the density of the demand shock $\mathbf{u}$, the marginal change in the probability of the event $\{\mathbf{q} \in T \cap S(0)\}$ is zero or even negative. Hence, nothing is gained by including vectors inside $S(0)$ in the trigger set.

It immediately follows from Proposition 1 that with optimal strategies, in this setting, price wars never occur in equilibrium. Related to this, the fact that a price war can only occur if a firm deviates implies that the price war probability is equal to the probability of detection of a deviation (whereas in Green-Porter, there are always type I and type II errors in the inference of deviant behaviour).

Now the analysis may look rather like that of the standard deterministic models: find the conditions under which collusion is sustainable (or find the maximal sustainable cartel price) given that firms use a simple grim trigger strategy. But this is still an imperfect-information environment in a fundamental sense. First, the fact that a price war can only occur if a firm in fact deviates is not a simplifying assumption but is derived as part of an optimal strategy, given the particular demand shocks used here. Second, and more importantly, while a price war can only occur if a firm deviates, this of course does not imply that any deviation will indeed be detected. Rather, marginal increases of the price war/detection probability (from zero) will have to suffice to deter marginal deviations.

3.3 The sustainability of collusion

We can now derive our main result, viz. the characterisation of how the sustainability of collusion depends on the degree of product differentiation. We have seen above that an optimal strategy is simply characterised by the cartel price p^c. A price war occurs if any vector $\mathbf{q} \in \Sigma - S(0)$ is observed, which leads to maximal punishment, i.e. breakdown of the cartel. Given that the demand functions are functions of the price difference k, the probability of a price war is a function of k as well. Denote this probability by $\alpha(k) = \text{Prob}\{\mathbf{q}(k) \in \Sigma - S(0)\}$, where $\alpha(0) = 0$.

Assuming that firm 2 adheres to the cartel price p^c, let us denote firm 1's expected per-period profit by $\pi_1(k, p^c)$. Moreover, let $\pi^c = \pi_1(0, p^c) = p^c/2$ and $\pi^b = p^b/2$ denote the collusive and Bertrand profits, respectively. Firm 1's expected discounted payoff from an infinitely repeated game is characterised by the Bellman equation

$$v_1(k, p^c) = \pi_1(k, p^c) + [1 - \alpha(k)]\delta v_1(k, p^c) + \alpha(k)\frac{\delta}{1-\delta}\pi^b$$

(cf. Green and Porter 1984). This equation can be be solved explicitly for v_1:

$$v_1(k, p^c) = \frac{\pi^b}{1-\delta} + \frac{\pi_1(k, p^c) - \pi^b}{1 - \delta + \delta\alpha(k)}. \tag{3}$$

For collusion to be sustainable, marginal deviations from p^c must not be profitable. By differentiating (3) with respect to k and considering the value $k = 0$ (where α vanishes), we obtain the condition

$$(1-\delta)\frac{\partial \pi_1(0, p^c)}{\partial k} \geq \delta(\pi^c - \pi^b)\frac{\partial \alpha(0)}{\partial k} \qquad \text{or}$$

$$\delta \;\geq\; \frac{\partial \pi_1(0, p^c)/\partial k}{(\partial \pi_1(0, p^c)/\partial k) + (\pi^c - \pi^b)(\partial \alpha(0)/\partial k)} =: \delta_0. \tag{4}$$

This constraint, of course, is familiar from Green and Porter (1984) and related works: the marginal gain of deviating from p^c must be counterbalanced by an expected loss in payoff due to an increase of the price war (= detection) probability α in order to deter deviations.

Where (4) is binding, v must be concave in k for $k = 0$ to be an optimum. If (4) holds with equality, the second-order condition is

$$(1-\delta)\frac{\partial^2 \pi_1(0, p^c)}{\partial k^2} \leq \delta(\pi^c - \pi^b)\frac{\partial^2 \alpha(0)}{\partial k^2}. \tag{5}$$

A first important result is that in this model, the sustainability of collusion as determined by (4) does not depend on the cartel price p^c. To see this, first notice that $\pi^c - \pi^b$ is simply $(p^c - p^b)/2$. The derivative $\partial \pi_1/\partial k$, evaluated at $k = 0$, is $1/2 +$

$p^c(\partial\tilde{q}/\partial k)$. On the other hand, p^b was determined by the equation $1/2+p^b(\partial\tilde{q}/\partial k)=0$. Therefore, we have $\partial\pi_1/\partial k = -(p^c - p^b)/(2p^b)$, and thus

$$\delta_0 = \frac{1}{1 - p^b(\partial\alpha/\partial k)}. \tag{6}$$

The important implication is that here, the first-order condition (4) relates to the sustainability of collusion *as such*, and not only with reference to a particular collusive price. This can be shown to be a general property of any model in which demand is a function of the *difference* of prices and is therefore a consequence of the Hotelling framework. In contrast, in some other models, an incentive constraint analogous to (4) relates to the sustainability of a particular price p^c. There, if for example the joint profit maximising price cannot be sustained, a lower collusive price might be sustainable.[14] A consequence for the analysis here is that in order to analyse the sustainability of of collusion, we need not determine the joint profit maximising price, which in this model would pose some problems.[15] As I will discuss further below, however, the analysis of the second-order condition (5) does lead to an upper bound for a sustainable cartel price.

The calculation of the price war probability α is relegated to Appendix B. There, it is shown that

$$\frac{\partial\alpha}{\partial k}(0) = \frac{1-\sigma}{\mu d p^b}\left(\log\frac{4}{2+e^{(1-\sigma)/\mu}+e^{-(1-\sigma)/\mu}}\right)^{-1}.$$

Combining the above results, we obtain

$$\delta_0 = \left[1 - \frac{1-\sigma}{\mu d}\left(\log\frac{4}{2+e^{(1-\sigma)/\mu}+e^{-(1-\sigma)/\mu}}\right)^{-1}\right]^{-1}. \tag{7}$$

[14] Cf. the representative-consumer model discussed in Raith (1996, p.108-121).

[15] In a standard Hotelling model, full market coverage is optimal provided the consumers' valuation for the good is sufficiently high, and then the joint profit maximising price is found by setting the utility of the worst-off consumer to zero (cf. Chang 1991). Here, this procedure does not work, since due to the unbounded support of the noise terms ε_{iz}, the probability of a purchase is positive, and less than unity, for every consumer and for any price.

Figure 2: Critical discount factor δ_0

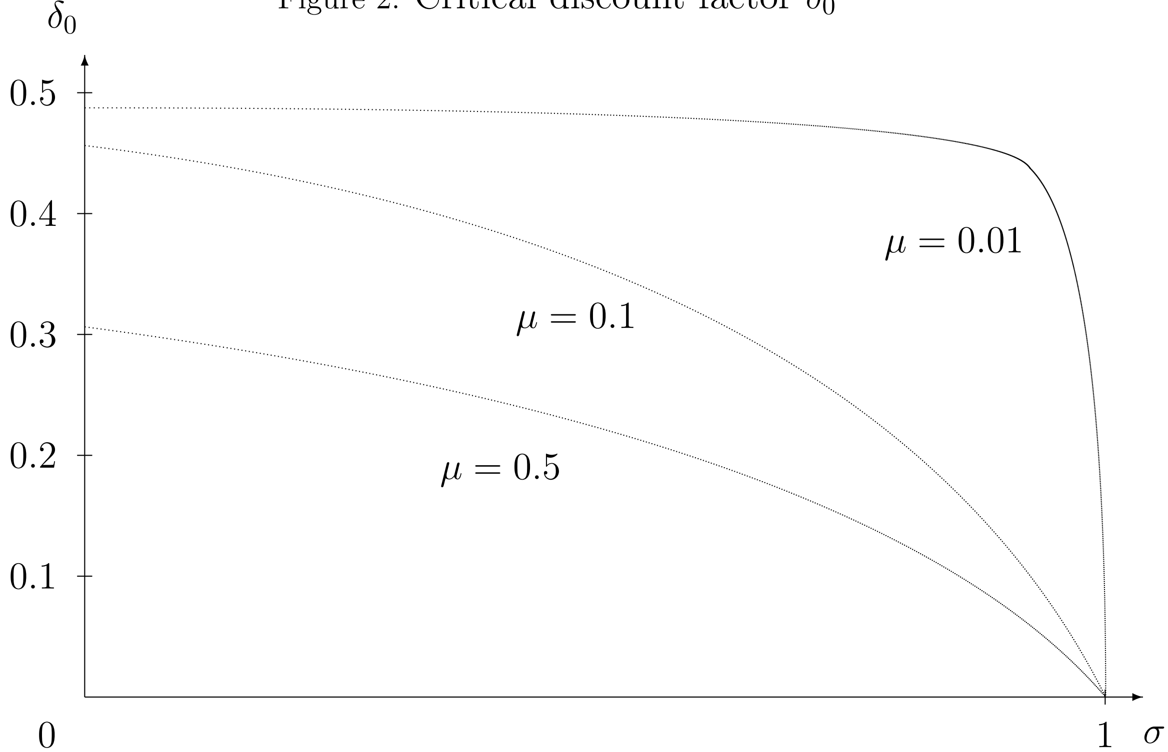

This expression is depicted in Figure 2 for different values of μ and for the case of maximal demand uncertainty, i.e. $d = 1$. The figure illustrates how δ_0 depends on σ and μ. This is stated formally in the following result.

Proposition 2 *(a) δ_0 is decreasing in σ and μ and increasing in d.*

(b) For $\sigma \to 1$ or $\mu \to \infty$, δ_0 converges to zero, and for $\mu \to 0$, δ_0 converges to $d/(1+d)$.

Proof: see Appendix B.

According to (4), collusion is sustainable if δ exceeds δ_0. Since δ_0 is decreasing in σ, we obtain the result that collusion is *less* sustainable for differentiated goods than for homogeneous goods. More precisely, while for larger discount factors collusion is always sustainable, for not too large discount factors there exists a critical level of differentiation such that for more homogeneous goods, collusion is sustainable, whereas it is not for more differentiated products. This result is consistent with the traditional view on the sustainability of collusion as a function of product differentiation, and states the opposite of what is often found in recent theoretical work.

To understand this relationship, consider the case of homogeneous goods. Here, the marginal gain from deviating may be rather large, but this is outweighed by an even larger (in absolute terms) increase of the detection probability. The reason for this is that, according to our assumption that the demand functions are highly correlated for equal prices, a deviation is relatively likely to lead to realisations of the demand vector outside the permissable region $S(0)$. This relative undesirability of marginal price cuts for large σ is reflected in a low value of δ_0.

Conversely, for more differentiated goods, the deterrent force of an increase of α vis-à-vis the increase of current-period profits is weaker, since a lower correlation of the demand functions corresponds to a larger set of demand vectors that occur with positive density. A price cut, therefore, is less likely to result in a realised demand vector outside this region.

The price and cross-price elasticities of demand, which of course vary with σ as well, play a double role, in such a way that in this model they do not affect the incentive constraint (4) at all: with a high elasticity in the case of homogeneous goods, marginal deviations may be tempting. But precisely a firm's ability for a firm to capture a large share of the market by cutting the price only by a small amount implies that such a deviation is likely to be detected. In this model, now, these two effects exactly cancel each other. This can be seen by considering the limit case $\mu \to 0$, which corresponds to the case that the demand functions are uncorrelated. Since according to Proposition 2, this limit does not depend on σ, this verifies that the negative relationship between δ_0 and σ indeed results from the correlation of the demand functions.

The comparative-statics properties of δ_0 with respect to μ and d are as follows: an decrease in μ means that the correlation of the demand functions decreases faster as products become more differentiated. This is reflected in an upward shift of $\delta_0(\sigma)$, i.e. a decrease of cartel stability. An increase an d has a similar effect, but for a different reason: while d does not affect the correlation of the demand shocks, it determines their variance. A more noisy environment is reflected in a decrease of α for any given σ and thus renders deviations more profitable, which is reflected in an upward shift of $\delta_0(\sigma)$.

As noted above, where the first-order condition is binding, the payoff function must be concave. Since $\pi_1(k, p^c)$ is concave, concavity of v would follow immediately if α were convex. This, however, is not the case: at $k = 0$, α is concave (see below). Thus, (5) places an upper bound on the sustainable price p^c, for which we can obtain a simple expression: with α as calculated in the Appendix it follows that

$$\frac{\partial^2 \alpha(0)}{\partial k^2} = \frac{(1-\sigma)^2}{4\mu^2 d^2 (p^b)^2} \frac{d^2 K - 2(1+d^2)}{K^2},$$

where $K = \log[4/(2 + \exp((1-\sigma)/\mu) + \exp(-(1-\sigma)/\mu))]$. Inserting this result and $\partial^2 \pi_1(0, p^c)/\partial k^2 = -1/p^b$ into (5) and eliminating K by rephrasing the first-order condition (4) as $K = -\delta(1-\sigma)/((1-\delta)\mu d)$, we obtain

$$p^c \leq p^b \left[1 + \frac{8\mu d}{2\mu(1-\delta)(1+d^2) + \delta d(1-\sigma)} \right]. \tag{8}$$

The sustainability of collusion can now be described more precisely as follows: for $\delta < \delta_0(\sigma = 0)$, a critical $\sigma(\delta)$ is uniquely defined by (7). For $\sigma < \sigma(\delta)$, collusion is not sustainable at any price because the first-order condition is violated. At $\sigma(\delta)$, any cartel price which does not exceed the upper bound given by (8) can be sustained. Since the payoff function is analytically intractable for arbitrary $k < 0$, and in particular since the optimal value of k for a deviating firm cannot be explicitly derived, it is not possible to determine analytically how the maximally sustainable price varies with σ for $\sigma > \sigma(\delta)$. We may conjecture, however, that there exists a monotonic relationship: if a price p^c is sustainable at some σ_0, then it will presumably be sustainable for any $\sigma > \sigma_0$, at least as long we restrict attention to price differences k for which $\alpha < 1$. If such a relationship holds, the maximally sustainable cartel price would be an increasing function of σ for $\sigma > \sigma(\delta)$. As noted above, however, this reasoning does not rule out the possibility that a firm might want to deviate by a large price cut, even if retaliation follows with certainty. This possibility is discussed in Appendix A.

4 Concluding remarks

I have argued in this paper that a satisfactory analysis of whether heterogeneity facilitates or hinders collusion should respect the importance of situations where firms undercut a collusive price only slightly (in the hope that this is not noticed). This requires a framework with uncertainty and imperfect monitoring. Moreover, uncertainty also provides an essential link between product heterogeneity and collusion: more heterogeneity leads to a decrease in the correlation of the firms' demand shocks. This implies an increase in uncertainty which in turn undermines the stability of collusion. This effect of product differentiation on the demand system is also likely to be a reason why firms producing heterogeneous products may find it difficult to reach a cartel agreement in the first place.

In his textbook, Tirole (1988, Chap. 5) describes product differentiation and collusion as two possible ways to escape the Bertrand paradox. The results of this paper, however, suggest that these two solutions are mutually exclusive: where the degree of differentiation is a choice parameter, firms can either raise margins by differentiating their products, or seek to abolish price competition, in which case products must be *standardised*. [16]

Economists have long believed that product heterogeneity is a factor hindering collusion among oligopolists, because it entails a situation of higher "complexity" than prevails with homogeneous goods. An analytical formulation of this argument, however,

[16] A quite different theory of why firms may want to standardise their products has been suggested by Klemperer (1992). In his model, customers demand an entire line of products and incur "shopping costs" for each of their suppliers. Firms have an incentive to offer standardised product lines such that each firm can satisfy a customer's demand. This relaxes price competition between the firms since, due to the switching costs, customers have an interest to maintain links with only one supplier. This theory is very elegant and convincing, but its application is restricted to markets in which customers demand a variety of products and incur sufficient switching costs (e.g., airlines' demand for aircraft). In contrast, the theory presented here seems applicable to a broader range of markets. While the two theories are complementary to some extent, it should be noted that the reasons for firms to standardise their products are very different in the two models: in Klemperer's model, firms standardise in order to relax noncooperative price competition, whereas here, standardisation serves to facilitate collusion.

has not been available. The purpose of this paper is to fill this gap in theory, i.e. to suggest a precise formulation of one kind of complexity that heterogeneity brings about. The conclusions of the present analysis suggest that it would be unwise to reverse the traditional view that heterogeneity hinders collusion by appealing to those recent theoretical models which rest on a deterministic structure.

A Cartel stability with large price cuts

A.1 A second incentive constraint

In this section, I will discuss the sustainability of collusion with respect to large price cuts. Sustainability of the cartel price p^c requires that v_1 as given by (3) has a global maximum at $k = 0$ (i.e. $p_1 = p_2 = p^c$), or

$$\frac{\pi^c - \pi^b}{1-\delta} \geq \frac{\pi_1(k, p^c) - \pi^b}{1 - \delta + \delta\alpha(k)} \quad \forall\, k,$$

which leads to the general incentive constraint

$$\frac{\pi_1(k, p^c) - \pi^b}{\pi^c - \pi^b} \leq \frac{\delta}{1-\delta}\alpha(k) \quad \forall k. \tag{9}$$

This general incentive constraint includes as a special case an incentive constraint which relates to marginal price cuts. This is the constraint (4), obtained by differentiation of v at $k = 0$, which has been emphasised in this paper and was analysed in Section 3.

Now, for not too large values of σ, the optimal value $k < 0$ which maximises the current-period profit of a deviating firm lies within the range where $\alpha(k) < 1$, since in the case of differentiated products, deviations are not detected very easily. For high values of σ, however, $\alpha(k) = 1$ is reached already for high (in absolute terms, small) values of k. In this case, the value of k which maximises $\pi(k, p^c)$, denoted $k^d(p^c)$, might be lower than the one where $\alpha(k) = 1$. Since in this range of k, the price war probability remains fixed at unity, the payoff $v_1(k, p^c)$ must be *decreasing* in k. Consequently, v has another local maximum at the optimal deviation price difference $k^d(p^c)$ in the region where $\alpha = 1$. This situation is depicted in Figure 3. A second constraint, therefore, relates to the case where a deviating firm sets $k = k^d(p^c)$ at a level where the deviation is detected with certainty. If we denote the resulting defection profit by π^d, the incentive constraint now becomes

$$\delta \geq \frac{\pi^d - \pi^c}{\pi^d - \pi^b} =: \delta_1. \tag{10}$$

This is the incentive constraint familiar from the theoretical literature discussed in the Introduction. For larger values of σ for which large discrete price cuts become relevant,

Figure 3: Payoff function $v_1(k, p^c)$

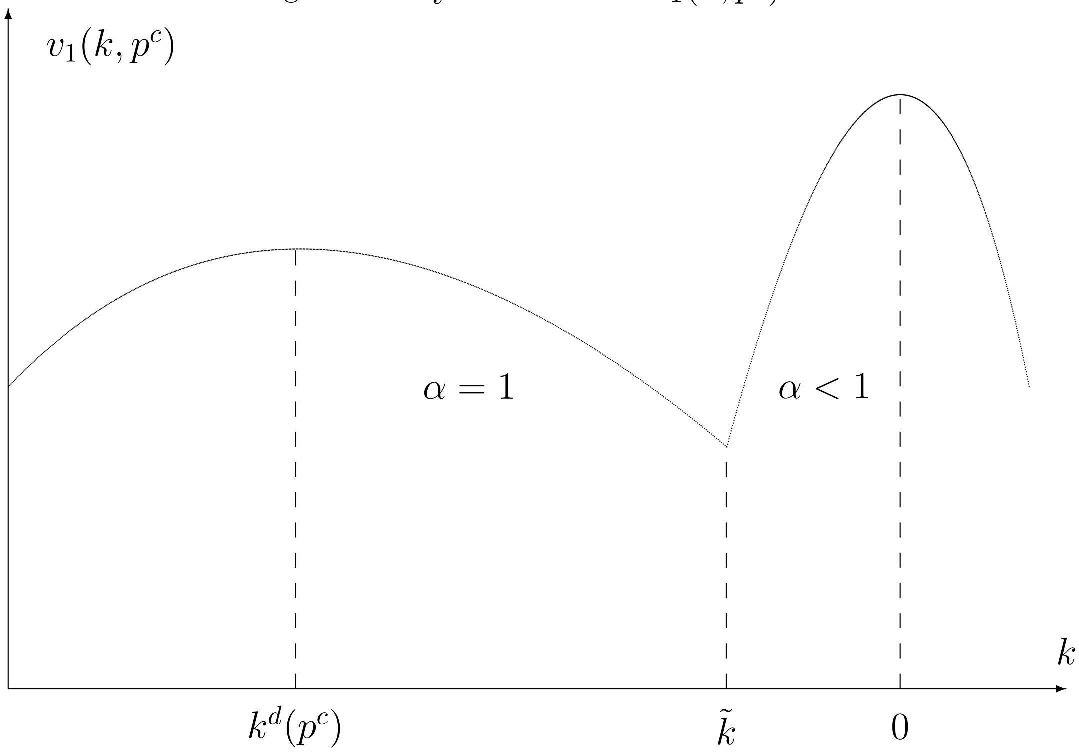

sustainability of a price p^c requires that both incentives constraints (4) and (10) are satisfied, i.e.

$$\delta \geq \max\{\delta_0, \delta_1\}.$$

Thus, the magnitude of δ_0 compared to δ_1 determines whether the sustainability of collusion is restricted by the threat of marginal deviations, as emphasised in this paper, or by large price cuts, as analysed in other theoretical papers. The values for δ_1 presented below should therefore be compared with the corresponding values of δ_0 depicted in Figure 2.

Clearly, an important factor determining the profitability of *marginal* price cuts, and hence the magnitude of δ_0, is the degree of uncertainty, parameterised by d. The larger d, the more important are marginal price cuts vis-á-vis large price cuts. Conversely, as d converges to zero, all uncertainty vanishes, any deviation from p^c is detected with certainty, and the sustainability of collusion is only determined by (10).

In the following, however, I will be concerned with the factors determining the profitability of *large* price cuts, and hence the magnitude of the critical discount factor

δ_1.

A.2 Evaluation of the critical discount factor δ_1

The optimal deviation profit π^d cannot be determined analytically. Most of what follows is therefore based on numerical simulations.

A first important observation is that δ_1 depends on p^c both absolutely and in comparison to the Bertrand price p^b. More precisely, the relative gain from deviating π^d/π^c is increasing in p^c (independent of p^b), while the benefit of collusion, π^c/π^b, equals p^c/p^b. As p^c approaches the Bertrand price, δ_1 approaches zero. This means that the incentive contraint (10) can be satisfied for any discount factor σ provided the cartel price is not too high above the Bertrand price. Conversely, as p^c approaches infinity, the deviation profit π^d converges to $2\pi^c$, and hence δ_1 converges to $1/2$. The sustainability of collusion as a function of product differentiation will therefore depend on how both p^c and the ratio p^c/p^b vary with σ. This question is dealt with in Section A.2.1.

In addition, capacity constraints limit the profitability of large price cuts and lower the critical value δ_1. This effect is illustrated in Section A.2.2. In the simulations presented below, I have fixed μ at the level 0.1.

A.2.1 Evaluation of δ_1 without capacity constraints

Constant cartel price: A natural starting point is to analyse how the critical discount factor δ_1 varies with σ if the cartel price is held constant. It turns out that δ_1 depends on σ in a way familiar from previous theoretical papers, as depicted in Fig. 4. That is, δ_1 is increasing in σ, implying that a given price p^c is *less* sustainable if products are more homogeneous. While π^c does not depend on σ, the Bertrand profit π^b is decreasing in σ, which, other things equal, would also imply a negative effect on δ_1. The deviation profit π^d, however, is increasing in σ. The shape of $\delta_1(\sigma)$ in Figure 4 reflects the fact that the latter effect dominates, which leads δ_1 to increase with σ. This means, as the goods become more homogeneous, the gain from deviating grows relatively faster than the benefit of colluding (or equivalently, the severity of punishment).

Figure 4: Critical discount factor δ_1 for constant p^c

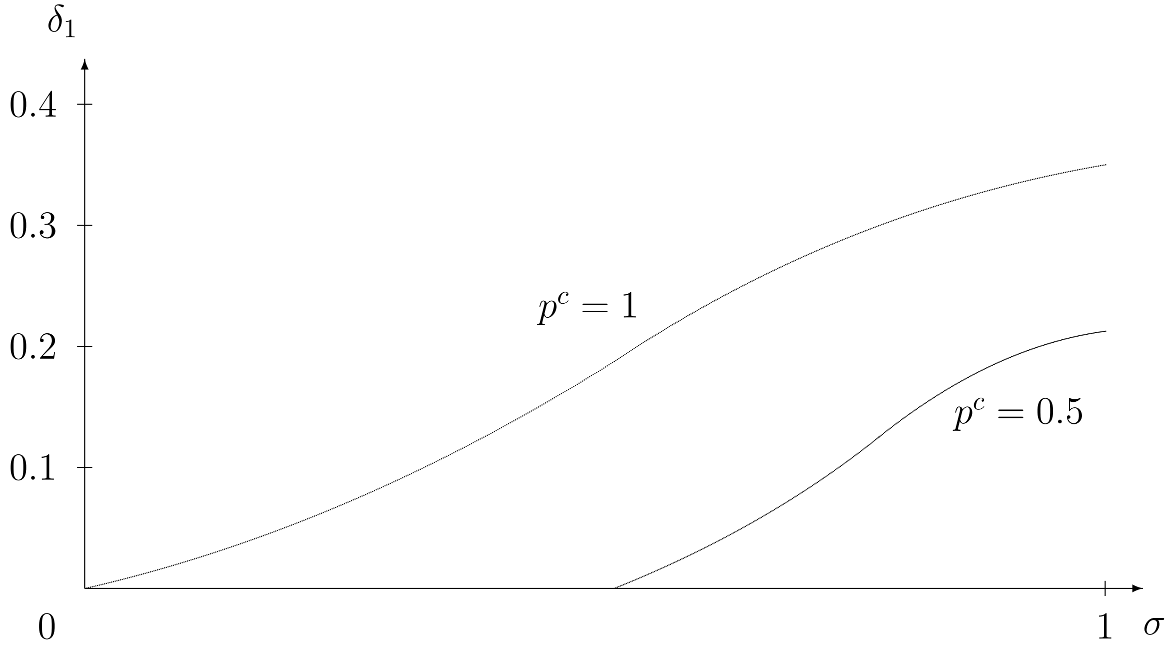

As mentioned above, a decrease in p^c shifts the schedule $\delta_1(\sigma)$ downwards (in the figure, the curve corresponding to $p^c = 1/2$ is only meaningful for sufficiently large values of σ where $p^c \geq p^b$).

Holding p^c constant implies that the ratio p^c/p^b increases with σ. This raises the question whether the fact that the difficulty in sustaining a certain price increases with the substitutablity of the goods, is to be attributed simply to the increase in the gap between the cartel price and the Bertrand price. To answer this question, therefore, we analyse how δ_1 depends on σ if p^c/p^b is held constant.

Cartel price as a multiple of the Bertrand price: If the cartel price to be sustained is a constant multiple of the Bertrand price (which implies that p^c, too, is decreasing in σ), a quite different picture compared to the previous case emerges.

It turns out that δ_1, calculated in this way, varies only little with σ and is in fact *decreasing* in σ. The reason is that not only the ratio π^c/π^b is constant (by construction), but also π^d is a multiple of π^b which varies only little with σ. This is due to two opposite effects which approximately cancel each other out: on one hand, for a given cartel price, an increase in σ leads to an increase of π^d. On the other hand, not only π^d itself, but also the ratio π^d/π^c is increasing in p^c; and since p^c is decreasing in σ, the same holds for π^d. Two examples: if $p^c = 2p^b$, then $\delta_1(\sigma)$ is around 0.2, and if $p^c = 1.5p^b$, then $\delta_1(\sigma)$ is approximately 0.11.

Sustainability of $p^{max}(\delta)$: In Section 3, we derived an upper bound for a sustainable cartel price for the case where the incentive constraint for marginal deviations (4) holds with equality, i.e. where $\sigma = \sigma(\delta)$. Moreover, we conjectured that this price (or even a higher price) is sustainable for more homogeneous goods, i.e. $\sigma > \sigma(\delta)$. The question arises whether at this threshold price, the second incentive constraint relating to large price cuts is satisfied.

To answer this question, we proceed as follows: for any given discount factor δ, calculate the corresponding critical degree of substitutability $\sigma(\delta)$ and then the upper bound for p^c according to (8). Using this value of p^c, we can compute δ_1 at $\sigma = 1$, where the maximum of $\delta_1(\sigma)$ is attained. Then, if δ_1 thus computed is below the actual δ we started with, the cartel price given by (8) can be sustained also with respect to large price cuts. The simulations suggest that this is indeed the case; in fact, for any given δ the corresponding δ_1 obtained by the procedure described above is always well below δ.

A.2.2 Capacity constraints

As argued in the Introduction, the possibility for deviating firms to capture a large share of the market in the short run is in practice limited by the presence of capacity constraints. Consequently, where such constraints play a role, the sustainability of collusion is largely determined by the incentive constraint relating to marginal deviations: a slight price cut which attracts some additional demand (that can still be satisfied) might be profitable, while a large price cut which leads to a price war for sure might not be considered a worthwhile option if the attainable short-run profit does not exceed the collusive profit by very much.

To get an idea how capacity constraints affect the two incentive constraints (4) and (10), consider the case $d = 1$ and assume that each firm has a capacity of one.

With this capacity, each firm can satisfy its random demand for any realisation of the density vector $\mathbf{u}$, *provided* both firms adhere to the cartel price, i.e. $k = 0$. This follows from $m_1 + m_2 = 1/2$ (cf. Section 2) and a maximum value of 2 for the densities u_1 and u_2. For a firm deviating marginally from p^c, the capacity constraint only has a second-order effect on expected profit and the price war probability. Hence, at this level of capacity, the incentive constraint for marginal deviations is not affected at all. For a firm deviating by a discrete amount, in contrast, the capacity constraint can be binding.

While it is straightforward but cumbersome to calculate precise expressions for a deviating firm's expected demand as a function of k, it is easy to establish that the expected demand can never exceed 5/6, as opposed to 1 without a capacity constraint: suppose $k = -\infty$ and therefore $q_1^l = q_1^r = 1/2$, which clearly maximises firm 1's demand. Then firm 1's demand is $(u_1 + u_2)/2$ if $u_1 + u_2 \leq 2$, and 1 otherwise, hence we have

$$E(q_1) = \int_{u_1 + u_2 \leq 2} \frac{u_1 + u_2}{2} f(\mathbf{u}) \mathrm{d}\mathbf{u} + \int_{u_1 + u_2 > 2} 1 \, f(\mathbf{u}) \mathrm{d}\mathbf{u} = \frac{1}{3} + \frac{1}{2} = \frac{5}{6}.$$

This lower level of expected demand compared to a situation without capacity constraints is reflected in a lower deviation payoff π^d and hence also in a lower critical discount factor δ_1. For example, while without capacity constraints δ_1 reaches a level of around 0.35 at $\sigma = 1$ and $p^c = 1$, it amounts to 0.26 with a capacity of 1. For $p^c = 2p^b$, δ_1 is below 0.14 rather that around 0.2, i.e. by a third lower.

These results confirm the assertion stated above: where capacity constraints play a role, the incentive constraint relating to marginal deviations becomes relatively more important vis-à-vis the constraint for large price cuts.

B Proofs

Proof of Proposition 1: 1. Consider a collusive strategy with trigger set T. Assuming that firm 2 adheres to the cartel price p^c, the expected discounted payoff for firm 1 is given by the Bellman equation

$$v_1(k,p^c) = \pi_1(k,p^c) + \delta v_1(k,p^c) \int\limits_{S(k)-T} f(k)\mathrm{d}\mathbf{q} + \delta v_1(k,p^c) \int\limits_{T} f(k)\eta(\mathbf{q})\mathrm{d}\mathbf{q}$$

(cf. Green and Porter (1984)), which can be rearranged to

$$v_1(k,p^c) = \pi(k,p^c)/(1-\delta w(k)) \quad \text{with}$$

$$w(k) = \int\limits_{S(k)-T} f(\mathbf{q},k)\mathrm{d}\mathbf{q} + \int\limits_{T} f(\mathbf{q},k)\eta(\mathbf{q})\mathrm{d}\mathbf{q}.$$

For collusion to be sustainable at p^c, we must have $\partial v_1(k,p^c)/\partial k = 0$ at $k=0$. Since $\partial \pi(k,p^c)/\partial k < 0$, $\partial w/\partial k$ must be positive for this condition to be satisfied. Since the inclusion of any vectors $\mathbf{q} \in \Sigma$ affects v_1 negatively, this can be worthwhile only if the inclusion results in an increase of $\partial w/\partial k$. We therefore want to examine how $w(k)$ changes with k.

2. We first show that $f(\mathbf{q},k)$ is constant for all $\mathbf{q} \in S(k)$ and therefore only depends on k, and that at $k=0$, we have $\partial f(\mathbf{q},k)/\partial k = 0$. The first observation follows immediately from the fact that $\mathbf{q}(k)$ is a one-to-one function of the uniformly distributed vector $\mathbf{u}$. From this it follows that $f(k)$ is simply the reciprocal of the Lebesgue measure of $S(k)$ in $\mathbf{q}$-space, denoted by $\lambda(S(k))$. To calculate this Lebesgue measure, let $\mathbf{u^A} = (1-d, 1-d)$, $\mathbf{u^B} = (1+d, 1-d)$, $\mathbf{u^C} = (1+d, 1+d)$ and $\mathbf{u^D} = (1-d, 1+d)$, and let $q_i^j = q_i(\mathbf{u^j})$ for $i = 1,2$ and $j = A,..,D$. Then the area of $S(k)$ is given by

$$\begin{aligned}
\lambda(S(k)) &= \frac{1}{2}\Big[(q_1^A - q_1^B)(q_2^A + q_2^B) + (q_1^B - q_1^C)(q_2^B + q_2^C) \\
&\quad + (q_1^C - q_1^D)(q_2^C + q_2^D) + (q_1^D - q_1^A)(q_2^D + q_2^A)\Big] \\
&= 4d^2(q_1^l q_2^r - q_1^r q_2^l),
\end{aligned}$$

and the derivative of this expression with respect to k at $k=0$ is 0.

3. Decompose T into the subsets $A = T - S(0)$ and $B = T \cap S(0)$, and decompose $S(k)$ into $S^A(k) = S(k) \cap (\Sigma - S(0))$ and $S^B(k) = S(k) \cap S(0)$. Then we have

$$\begin{aligned} w(k) &= \int_{S^A(k)-A} f(k)\mathrm{d}\mathbf{q} + \int_{S^B(k)-B} f(k)\mathrm{d}\mathbf{q} \\ &+ \int_{S^A(k)\cap A} f(k)\eta(\mathbf{q})\mathrm{d}\mathbf{q} + \int_{S^B(k)\cap B} f(k)\eta(\mathbf{q})\mathrm{d}\mathbf{q}. \end{aligned}$$

Since

$$\frac{\partial\lambda}{\partial k}[S^A(k) - A] = -\frac{\partial\lambda}{\partial k}[S^A(k) \cap A] > 0,$$

it follows from argument (2) above and $\eta(\mathbf{q}) < 1$ that

$$\frac{\partial}{\partial k}\left[\int_{S^A(k)-A} f(k)\mathrm{d}\mathbf{q} + \int_{S^A(k)\cap A} f(k)\eta(\mathbf{q})\mathrm{d}\mathbf{q}\right] > 0.$$

Thus, since vectors $\mathbf{q} \in A$ do by construction not affect the payoff v_1 in equilibrium but affect $\partial v_1/\partial k$ positively, it is optimal to set $A = \Sigma - S(0)$. Similarly, since

$$\frac{\partial\lambda}{\partial k}[S^B(k) - B] = -\frac{\partial\lambda}{\partial k}[S^B(k) \cap B] < 0,$$

it follows that

$$\frac{\partial}{\partial k}\left[\int_{S^B(k)-B} f(k)\mathrm{d}\mathbf{q} + \int_{S^B(k)\cap B} f(k)\eta(\mathbf{q})\mathrm{d}\mathbf{q}\right] < 0.$$

Thus, since vectors $\mathbf{q} \in B$ by construction affect v_1 negatively but also affect $\partial v_1/\partial k$ negatively, it is optimal to set $B = \emptyset$. ■

Calculation of the price war probability $\alpha(k)$: 1. If $k = 0$, we have $q_1 = m_1u_1 + m_2u_2$ and $q_2 = m_2u_1 + m_1u_2$, and given the distribution of $\mathbf{u}$, the support $S(0)$ of $\mathbf{q}(0)$ is a parallelogram enclosed by the lines corresponding to $u_1 = 1 - d$, $u_1 = 1 + d$, $u_2 = 1 - d$ and $u_2 = 1 + d$. For any fixed u_2^0, the points $\mathbf{q}$ on the line $u_2 = u_2^0$ are characterised by

$$q_1/m_2 - q_2/m_1 = (m_1/m_2 - m_2/m_1)u_1,$$

which is positive because $m_1 > m_2$. If firm 1 undercuts the cartel price, the difference $q_1/m_2 - q_2/m_1$ can only increase. Therefore, firm 1 has surely deviated if

$$q_1/m_2 - q_2/m_1 > (m_1/m_2 - m_2/m_1)(1+d). \tag{11}$$

Similarly, for $k = 0$ we have

$$q_1/m_1 - q_2/m_2 = (m_2/m_1 - m_1/m_2)u_2,$$

which is negative, and firm 1 has surely deviated if

$$q_1/m_2 - q_2/m_1 > (m_2/m_1 - m_1/m_2)(1-d). \tag{12}$$

2. For $k \leq 0$, condition (11) becomes

$$(m_1 q_1^l - m_2 q_2^l)u_1 + (m_1 q_1^r - m_2 q_2^r)u_2 > (m_1^2 - m_2^2)(1+d), \tag{13}$$

and condition (12) becomes

$$(m_2 q_1^l - m_1 q_2^l)u_1 - (m_1 q_2^r - m_2 q_1^r)u_2 > -(m_1^2 - m_2^2)(1-d). \tag{14}$$

These inequalities describe downward-sloping and upward-sloping lines in **u**-space, respectively, as shown in Figure 5. It is straightforward to establish that the intercepts of (13), $\underline{u}_1$ and $\bar{u}_1$, indeed lie in the interval $[1-d, 1+d]$, as depicted in the figure. The two curves intersect at

$$\tilde{u}_1 = \frac{m_2 q_1^r - m_1 q_2^r}{q_1^r q_2^l - q_1^l q_2^r}(1+d) + \frac{m_1 q_1^r - m_2 q_2^r}{q_1^r q_2^l - q_1^l q_2^r}(1-d).$$

The price war probability α is given by the area to the right of line (13) joined by the area below line (14), and divided by the density $1/(4d^2)$. Denote by $u_2^a(u_1)$ and $u_2^b(u_1)$ the values of u_2 as a function of u_1 given by (13) and (14), respectively. Then we have

$$\begin{aligned} \alpha &= \frac{1}{4d^2}\left\{ \int_{1-d}^{\underline{u}_1} [u_2^b(u_1) - (1-d)]\mathrm{d}u_1 + \int_{\underline{u}_1}^{\tilde{u}_1} [(1+d-u_2^a(u_1)) + \right. \\ &+ \left. (u_2^b(u_1) - (1-d))]\mathrm{d}u_1 + 2d(1+d-\tilde{u}_1) \right\} \end{aligned}$$

Substitution of q_1^l, q_1^r, q_2^l, and q_2^r into this expression and differentiation with respect to k at $k = 0$ leads to expressions for $\partial\alpha/\partial k$ and $\partial^2\alpha/\partial k^2$ stated in the text. ■

Figure 5: Calculation of $\alpha(k)$

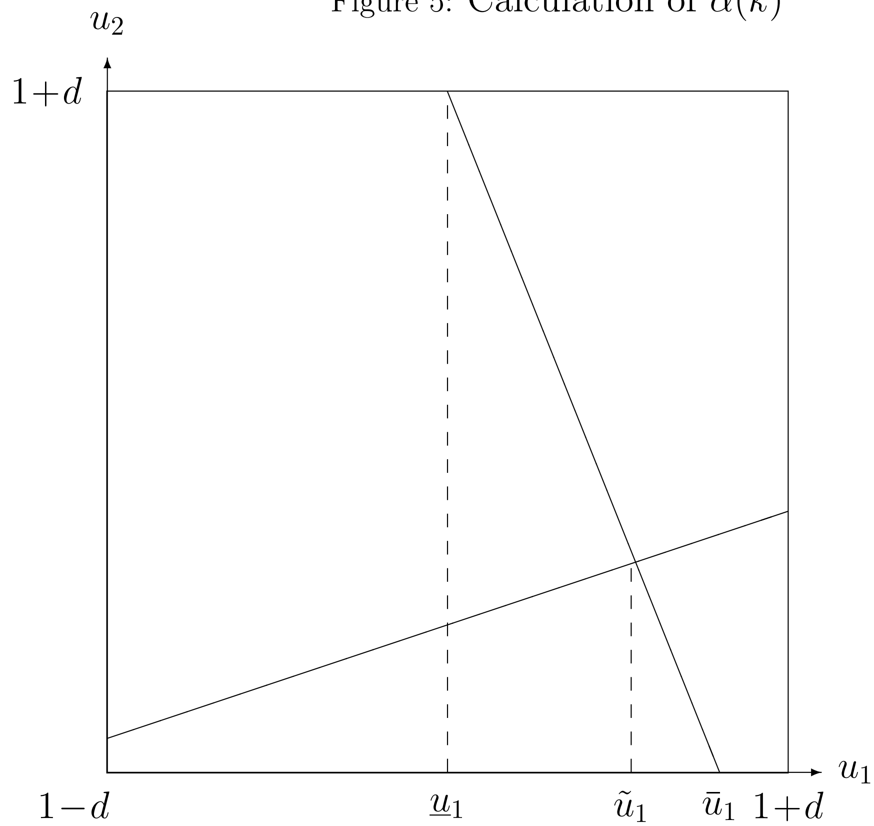

Proof of Proposition 2: Part (a): The results are obtained by writing δ_0 as

$$\delta_0 = \left[1 - \frac{z}{d}\left(\log \frac{4}{2 + e^z + e^{-z}}\right)\right]^{-1} \tag{15}$$

with $z = (1 - \sigma)/\mu$. The derivative with respect to d is obvious. The derivative with respect to z has the same sign as

$$(e^z - 1)z + (e^z + 1)\log \frac{4}{2 + e^z + e^{-z}}. \tag{16}$$

This expression is 0 at $z = 0$, and the derivative with respect to z is

$$e^z(z + \log \frac{4}{2 + e^z + e^{-z}}), \tag{17}$$

which again is 0 at $z = 0$ and increasing in z. Since (17) is therefore positive, it follows that (16) is positive. Hence, δ_0 is increasing in z, and therefore decreasing in both σ and μ.

Part (b): The limits of δ_0 for $\sigma \to 1$ and $\mu \to \infty$ are obtained by taking the limit $z \to 0$ (using l'Hôpital's rule) in (15), and the limit $\mu \to 0$ by taking the limit $z \to 0$.

■

References

Abreu, Dilip 1986: Extremal equilibria of oligopolistic supergames, *Journal of Economic Theory* 39: 191-225

Abreu, Dilip; David Pearce and Ennio Stacchetti 1986: Optimal cartel equilibria with imperfect monitoring, *Journal of Economic Theory* 39: 251-269

Archibald, G.C; B.C. Eaton and R.G. Lipsey 1986: Address models of value theory, in: J.E. Stiglitz and G.F. Mathewson (eds.), *New developments in the analysis of market structure* (MIT), 3-47

Chang, Myong-Hun 1991: The effects of product differentiation on collusive pricing; *International Journal of Industrial Organization* 9: 453-469

Clarke, Roger 1985: Industrial Economics (Basil Blackwell)

d'Aspremont, Claude; Jean Jaskold Gabszewicz and Jacques-François Thisse 1979: On Hotelling's 'Stability in Competition', *Econometrica* 47: 1145-50

Deneckere, Raymond 1983: Duopoly supergames with product differentiation, *Economics Letters* 11: 37-42

de Palma, André; Robin Lindsey; Balder von Hohenbalken; Douglas S. West 1994: Spatial price and variety competition in an urban retail market - A nested logit analysis, *International Journal of Industrial Organization* 12: 331-357

de Palma, André; V. Ginsburgh, Y. Y. Papageorgiou, J.-F. Thisse 1985: The principle of minimal differentiation holds under sufficient heterogeneity, *Econometrica* 53: 767-781

Fraas, Arthur G.; and Douglas Greer 1977: Market structure and price collusion: An empirical analysis, *Journal of Industrial Economics* 26: 21-44

Friedman, James W.; and Jacques-François Thisse 1993: Partial collusion yields minimal product differentiation, *RAND Journal of Economics* 24: 631-645

Fudenberg, Drew; and David Levine 1991: An approximate folk theorem with imperfect private information, *Journal of Economic Theory* 54:26-47

Green, Edward J.; and Robert H. Porter 1984: Noncooperative collusion under imperfect price information, *Econometrica* 52: 87-100

Häckner, Jonas 1993: Product differentiation and the sustainability of collusion, in: Häckner, Price and quality - essays on product differentiation, Stockholm (IUI Dissertation Series)

Häckner, Jonas 1994: Collusive pricing in markets for vertically differentiated products, *International Journal of Industrial Organization* 12: 155-177

Hay, George A.; and Daniel Kelley: An empirical survey of price-fixing conspiracies, *Journal of Law and Economics* 17: 13-38

Jehiel, Philippe 1992: Product differentiation and price collusion; *International Journal of Industrial Organization*

Lehrer, Ehud 1992: On the equilibrium payoffs set of two player repeated games with imperfect monitoring, *International Journal of Game Theory* 20: 211-226

Martin, Stephen 1993: Advanced Industrial Economics (Blackwell)

Porter, Robert 1983: Optimal cartel trigger-price strategies, *Journal of Economic Theory* 29: 313-338

Posner, Richard A. 1976: Antitrust Law - An Economic Perspective, Chicago and London (Univ. of Chicago Press)

Raith, Michael 1996: Product differentiation, uncertainty and price coordination in oligopoly, unpublished Ph.D. thesis, University of London

Ross, Thomas W. 1992: Cartel stability and product differentiation, *International Journal of Industrial Organization* 10:1-13

Scherer, F.M.; and David Ross 1990: Industrial market structure and economic performance (Houghton Mifflin)

Stigler, George 1964: A theory of oligopoly, *Journal of Political Economy* 72:44-61

Sultan, Ralph G.M. 1974: Pricing in the Electrical Oligopoly, Vol. I, Cambridge/Mass. (Harvard)

Tirole, Jean 1988: The Theory of Industrial Organization, Cambridge (Mass.) and London (MIT)

Ulen, Thomas 1983: Railroad cartels before 1887: the effectiveness of private enforcement of collusion, *Research in Economic History* 8: 125-144

Verboven, Frank 1995: Localized competition, multimarket operation and collusive behavior, mimeo, CentER Tilburg

Wernerfelt, Birger 1989: Tacit collusion in differentiated Cournot games, *Economics Letters* 29: 303-306

Zhang, Z. John 1995: Price-matching policy and the principle of minimal differentiation, *Journal of Industrial Economics* 43: 287-299