The Determinants of Collusion under Exogenous Demand Fluctuations

Erik Pot, [1] Ronald Peeters, [2] Hans Peters, [3] and Dries Vermeulen [4]

Abstract

In this paper attempt to reconcile the–at first sight different–views on the determinants of collusion and price wars expressed in Rotemberg and Saloner (1986), Green and Porter (1984), and Stigler (1964). We first argue that the logic of R&S presupposes two determinants for collusion, namely (1) market shares are publicly observable, and (2) volatility of market shares due to exogenous factors is limited. We make our arguments in a model in which firms repeatedly play a Bertrand type price competition game. Following R&S we show under the two conditions of public observability and limited volatility of market shares that within the model firms can collude using dynamic price adjustment strategies. We show that when the first condition (public observability) is violated, we revert to the logic of Green and Porter. When the second condition (limited volatility of market shares) is violated, for example when consumer loyalty has decreased, we also observe that collusion can no longer be sustained, in line with the arguments in Stigler (1964).

Keywords and Phrases: Noncooperative collusion, price wars, repeated games.

JEL Classification Numbers: C73, D43, L13.

[1] e.pot@maastrichtuniversity.nl. Department of Quantitative Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands.

[2] r.peeters@maastrichtuniversity.nl. Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands.

[3] h.peters@maastrichtuniversity.nl. Department of Quantitative Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands.

[4] d.vermeulen@maastrichtuniversity.nl. Department of Quantitative Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands.

1 Introduction

This paper is a contribution to the ongoing discussion on the stability of collusion and the conditions under which pricing agreements among oligopolists can be sustained in equilibrium.

It is well known in the industrial organization literature that fluctuations in levels of both individual and market demand play an important role in the stability of collusion and the occurrence of price wars. An early contribution in this context is Stigler (1964). In the model of Stigler firms face uncertainty regarding their individual demand, and they cannot directly observe their opponents' behavior. An unexpectedly large drop in a firm's own individual demand may therefore be attributed to an (unobserved) deviation from collusion by one of the opponents, and a price war is seen as the reversion to competitive behavior to punish such deviations. Stigler argued that, in a market with high consumer loyalty, deviations from a collusive agreement are relatively easy to detect. Therefore collusion is easier to sustain in markets with high consumer loyalty than in markets with low consumer loyalty.

The first equilibrium based paper showing how price wars can occur on the equilibrium path is Green and Porter (1984). In their model, a firm can experience a period of unexpectedly bad performance both as a result of deviating behavior by one of the firms and as a result of (unobservable) low aggregate demand. Since, because of unobservability, firms are unable to distinguish between these two scenarios, they have to revert to retaliatory behavior in either case in order to discourage deviating behavior. Price wars will thus occur with certainty in periods of low individual demand, even when deviations did not occur.

Another milestone in the discussion is the paper by Rotemberg and Saloner (1986). In a model with volatile aggregate demand and fixed market shares they show that partial collusion can be sustained in equilibrium using countercyclical pricing strategies. The logic of their argument is that, during a boom, the temptation for firms to deviate from the collusive agreement to attract consumers starts to outweigh the decrease in profits resulting from the ensuing retaliatory price war. To counterbalance this threat to collusion, in equilibrium firms consequently employ a gradual and coordinated downward adjustment of the price levels in response to the increased level of demand in periods where the market is booming. Since the decision to decrease prices during a boom is taken jointly by all competitors as part of the collusive agreement such an orchestrated and voluntary decrease in prices is in fact not a price war, but can better be viewed as a form of dynamic collusion where prices are deliberately adjusted to the circumstances, precisely with the intention to stabilize collusion. A full blown price war in the sense of full reversion to marginal cost pricing does not occur on the equilibrium path. This was concisely put by Ellison (1994): "Rotemberg and Saloner (1986), is commonly associated with the statement that price wars are more likely to occur during booms, and therefore viewed as somehow in opposition to the Green and Porter theory. The actual Rotemberg and Saloner model, however, is really about countercyclical pricing – firms have perfect information and adjust prices smoothly in response to demand conditions."

In this paper we present a model where firms interact repeatedly in a Bertrand type model based on price competition. Each firm can choose to collude, to price at marginal cost, or to deviate from collusion (undercutting the collusive price). We assume that aggregate demand is inelastic and that the division of market shares may fluctuate over time. Fluctuations of market shares model exogenous factors such as for example consumer loyalty (or better: lack thereof) and location effects. Typically such factors are outside the control of the firms and are known to affect collusive opportunities (Stigler (1964), Green and Porter (1984)).

Within this model, both with private information on individual market shares and with public information, we derive the conditions under which strategy profiles in trigger strategies where each firm chooses to collude unless a deviation has been detected in the past (in which case firms revert to marginal cost pricing) can be sustained as a perfect Bayesian Nash equilibrium.

In both the case of private and the case of public information we find that collusive behavior in trigger strategies is harder to sustain when market shares have high volatility over time, and periods with low individual demand are possible. Moreover, in the public information case, opportunities for partial collusion enhance collusion. The logic driving these results is fairly intuitive. Collusive behavior can be sustained in equilibrium by trigger strategies precisely when expected profits for firms adhering to the collusive agreement are higher than the single period gains from deviation. This condition is particularly stringent for periods where individual demand is low, since in such periods the expected profits when a firm follows the agreement are minimal, while the immediate gains from deviation (undercutting) are high. In addition, when a firm also expects its individual demand to be low in the future, which is more likely when current individual demand is low, the punishment ensuing the breaking of the agreement is relatively small. Hence, high volatility of individual market shares hampers collusion. On the other hand, when partial collusion is possible, and market shares are observable, a countercyclical pricing policy with partial collusion as in Rotemberg and Saloner can be applied to enhance sustainability of collusive behaviour in equilibrium.

These results can be seen as an attempt to reconcile the views expressed in Stigler (1964) and Green and Porter (1984) with the arguments in Rotemberg and Saloner (1986). In our model with private information, increased volatility of market shares prevents collusion. This is in line with the view of Green and Porter, who argue that collusion is most likely to break down in periods of low demand, and the view of Stigler, who argued that high consumer loyalty, which is directly related to low volatility of market shares, is one of the stabilizing factors for collusion.

On the other hand, in our model with public observability of market shares, the basic logic of the arguments in Rotemberg and Saloner is still valid. Firms easily collude when differences in individual demand remain relatively small. [1] However, in the presence of fluctuations of market shares collusion becomes more difficult to sustain. The cartel is then stabilized in periods of high fluctuations in market shares by using a coordinated price adjustment scheme. When market shares are out of balance, a policy of lower collusive price setting discourages deviations by firms with lower market shares.

Thus, our argument here is that the driving force behind the results of Rotemberg and Saloner is not the presence of shocks on total demand per se, but more in general the public observability of market shares in conjunction with low volatility of these market shares. Public observability of market shares is essential to the implementation of dynamic price strategies and hence partial collusion, while low volatility of market shares guarantees that dynamic adjustment of prices via a collusive agreement can sufficiently decrease gains from deviation.

Note that both conditions, public observability and low volatility, are fulfilled in Rotemberg and Saloner. They assume that total demand is publicly observed, and moreover that total demand is always equally divided over firms, so that implicitly the division of market shares over firms is common knowledge. Thus, firms are assumed to have full information on market shares. And indeed, in an environment where firms can make binding agreements on market shares, the full force of Rotemberg and Saloner's arguments applies, and collusion, at least partial collusion, can be sustained in equilibrium using countercyclical pricing strategies.

However, we show that in addition to this basic observation, the logic of Rotemberg and Saloner breaks down as soon as one of the two conditions is violated. When in our model market shares are no longer publicly observable, collusion via dynamic price schemes is no longer possible,

[1] Despite the differences between the R&S model and our model, the intuition is the same in both models. The central issue is changes in potential gains from deviation. In the model of R&S these changes are the result of changes in aggregate demand under fixed market shares. In our model these changes are conversely the result of changes in *individual* demand (market shares) under *constant* aggregate demand. Nevertheless the effect is the same: both in the case of high aggregate demand with fixed market shares and in the case of constant aggregate demand with low market shares gains from deviation are increased.

and we effectively revert to a model where the logic of Green and Porter applies. Hence, the conclusions of Green and Porter versus Rotemberg and Saloner are not contradictory, but rather complement each other, and can be observed under different conditions within a single dynamic model of Bertrand competition.

Also, when market shares are still observable but the volatility of market shares is sufficiently high, dynamic pricing schemes can no longer be sustained in equilibrium. This is due to the fact that large changes in market shares can no longer be compensated for by a countercyclical collusive price adjustment scheme. The gains from deviation for firms with low individual demand simply can no longer be counterbalanced in that case. Then, in line with the findings of Stigler, collusion breaks down.

In conclusion, the logic of Rotemberg and Saloner not only presupposes observability of market shares; also low volatility of individual demand, for example via a sufficient amount of control over individual demand, is essential to their argument. So, mere observability is not sufficient. When market shares cannot be enforced collectively and are subject to large exogenous fluctuations that are outside the control of the firms, the conclusions of Rotemberg and Saloner are mitigated by the volatility of market shares, and collusion will be harder, or even impossible, to sustain.

2 Basic model

In this section we present the basic model we use in our analysis. In our model firms interact repeatedly in a Bertrand-type pricing game. We first define the one-shot game played between firms in each period of time. In each period firms simultaneously and independently choose one out of three pricing strategies, namely collusion, marginal cost pricing, and undercutting. Based on the choice of strategies of the firms each firm receives a payoff. This basic game is then repeated over an infinite time horizon. Firms use discounting to evaluate the resulting payoff streams.

2.1 The one shot game

The one shot game is based on a Bertrand model in which n firms compete on price in a market for a homogeneous good. In the one shot game each firm has three available actions it could possibly take, namely to collude (C), to undercut (U), or to price at marginal cost (M). Thus, each firm chooses an action $a_i \in \{U, C, M\}$ without prior information on the choices of the other firms. A second ingredient of the model is a vector $\varphi = (\varphi_1, \ldots, \varphi_n)$ of market shares, where $\varphi_i \geq 0$ represents the market share of firm i. The market share vector prescribes how total profits are divided among firms when all firms choose the same action. Market shares divide the total market, so $\sum_i \varphi_i = 1$.

To be precise, profits are determined as follows. Given a profile $a = (a_i)_{i \in N}$ of chosen actions, the resulting profit of firm i is denoted by $\Pi_i(a)$. When $a_i = C$ for all i, then $\Pi_i(a) = \varphi_i \Pi$, where Π represents the monopoly profit in the market. Thus, in this case, firms act monopolistically and profits are divided according to market share. When there is a firm k with $a_k = U$ and $a_i = C$ for all $i \neq k$, then $\Pi_k(a) = \Pi$ and $\Pi_i(a) = 0$ for all $i \neq k$. In all other cases all profits are zero.

Although our model is based on the model of Bertrand competition, it obviously deviates from the standard approach in modeling Bertrand competition in two aspects. First, our approach starts with the observation that, even though in the original model each firm can choose any price, only three strategies for price setting are relevant for the dynamics of the Bertrand model. These three choices are (a) collusive pricing, (b) price setting slightly below collusive pricing to exploit collusive behavior and to capture the market, and (c) competitive (marginal cost) pricing. In our approach we discard other possible choices and only focus on these three crucial price setting strategies. That way we try to keep the analysis of the one shot game simple while we preserve the essential ingredients of the Bertrand model.

The second aspect regards our choice of the payoffs in the one shot model. It is clear that profits are zero when at least one of the firms chooses to use marginal cost pricing. The firm that uses this action attracts the market, but does not make any profit, while the other firms do not have any customers. In the case that more than one firm undercuts the collusive price, one might assume that the undercutting firms take their share of the monopoly profit. However, for simplicity we adopt the more radical assumption that also in this case all profits are zero.

The one shot game has several Nash equilibria. Evidently the action profile in which all firms play M is a Nash equilibrium. However, when there are at least three firms, an action profile in which for example one firm plays C and all other firms play M is also a Nash equilibrium. Nevertheless, the two important observations to make here are (1) that the action profile in which all firms play C, the collusion strategy profile, is not a Nash equilibrium, and (2) that in any symmetric Nash equilibrium (either all firms playing M or all firms playing U) of the one shot game all firms receive zero profit.

2.2 The repeated game

In the repeated game the one shot game is repeated over an infinite time horizon. At the start of each period $t = 0, 1, 2, \ldots$ the vector $\varphi_t = (\varphi_{1t}, \ldots, \varphi_{nt})$ of market shares is determined. This vector is stochastic, and it could in general depend on both the actions taken previously by the firms and on the realizations of market shares in earlier periods. Nevertheless, for simplicity we only study exogenous processes where the realizations do not depend on the actions taken by the firms.

Next, when φ_t is realized, each firm receives information h_{it}. Typically h_{it} is a record of all actions taken by firms in earlier periods, the realized market shares of all firms in earlier periods, and either a firm's own market share in the current period (*private information*), or all realized market shares in the current period (*public information*).

A strategy for firm i in the repeated game is a function s_i that prescribes for each history h_{it} the action $s_i(h_{it}) \in \{U, C, M\}$ that firm i chooses. We require $s_i(h_{it})$ to be specified for any conceivable history h_{it}, not just for those that are actually realized by previous actions of the firms. This is standard practice in game theoretic models and facilitates defining the concept of (subgame perfect) Nash equilibrium. In particular, a firm should not only specify what it will do when all firms act according to agreement, but also how it will react to conceivable deviations from the agreement (and by extension also to deviations from deviations from the agreement, etc).

We write $s(h_t) = (s_1(h_{1t}), \ldots, s_n(h_{nt}))$ for the profile of actions that is played at time t given the information $h_t = (h_{1t}, \ldots, h_{nt})$. Let s^t denote the map $h_t \mapsto s(h_t)$. The initial division of market shares is given by $\varphi_0 = (\varphi_{10}, \ldots, \varphi_{n0})$. Further, the associated information to the firms is denoted by $h_0 = (h_{10}, \ldots, h_{n0})$. The density function $f_{it+1}(\varphi_{it+1} \mid \varphi_{it})$ denotes the density of the probability distribution of φ_{it+1} conditional on the event that at time t the market share of firm i is φ_{it}. By $\mathbb{E}(\Pi_i(s^t) \mid h_{i0})$ we denote the expected value of the profit to firm i at time t, given the strategy profile s and the initial information h_{i0} of firm i.

Given the profile $s = (s_1, \ldots, s_n)$ of strategies, firm i evaluates the resulting stream

$$\mathbb{E}(\Pi_i(s^0) \mid h_{i0}),\ \mathbb{E}(\Pi_i(s^1) \mid h_{i0}),\ \ldots$$

of expected profits via the discounted criterion defined by

$$\Pi_i(s \mid h_{i0}) = \sum_{t=0}^{\infty} \delta^t \cdot \mathbb{E}(\Pi_i(s^t) \mid h_{i0}).$$

Given a strategy profile s and a strategy s'_i for firm i, let (s, s'_i) denote the strategy profile where all firms $j \neq i$ play according to the strategy s_j, while firm i plays according to strategy s'_i. A strategy profile s is a Bayesian Nash equilibrium when for every firm i

$$\Pi_i(s \mid h_{i0}) \geq \Pi_i((s, s'_i) \mid h_{i0})$$

holds for every strategy s'_i of firm i and any initial information h_{i0}.

Analogously, let $\mathbb{E}(\Pi_i(s^{t+k}) \mid h_{it})$ denote the present value of the profit of firm i at time $t+k$, given the strategy profile s and information h_{it} to firm i at time t. Write

$$\Pi_i(s \mid h_{it}) = \sum_{k=0}^{\infty} \delta^k \cdot \mathbb{E}(\Pi_i(s^{t+k}) \mid h_{it})$$

for the present value of the stream of profits to firm i at information set h_{it}. A Nash equilibrium s is a perfect Bayesian Nash equilibrium when, at every information set h_{it},

$$\Pi_i(s \mid h_{it}) \geq \Pi_i((s, s'_i) \mid h_{it})$$

holds for every strategy s'_i of firm i.

3 Collusive equilibria under private information

In this section we assume private information. Thus, for every i and t, h_{it} consists of φ_{it} together with all realizations of market shares and all actions taken by the firms in all previous rounds. We also assume that all realizations of market shares φ_{it} are within an interval $[\underline{\varphi}, \overline{\varphi}]$ with $0 \leq \underline{\varphi} < \overline{\varphi} \leq 1$. We first analyze under what conditions collusion can be sustained as a perfect Bayesian Nash equilibrium via trigger strategies.

TRIGGER STRATEGIES The trigger strategy T_i of firm i is defined by

$$T_i(h_{it}) = \begin{cases} C & \text{if all firms chose action } C \text{ in all previous rounds according to } h_{it} \\ M & \text{otherwise} \end{cases}$$

and $T = (T_i)_{i \in N}$ denotes the profile of trigger strategies. We derive the following necessary and sufficient condition for T to be a perfect Bayesian Nash equilibrium.

Theorem 3.1 *The strategy profile T is a perfect Bayesian Nash equilibrium precisely when*

$$\sum_{k=0}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \geq 1$$

holds for any possible market share $\varphi_{it} \in [\underline{\varphi}, \overline{\varphi}]$ at time t, for all firms i and at every time t.

Proof. Due to the one deviation property (see e.g. Hendon et al. (1996)), the trigger strategy profile is a perfect Bayesian Nash equilibrium exactly when every firm i, at every time t, and at every information set h_{it} the trigger strategy renders at least the same expected profit as an instantaneous deviation. Thus, consider firm i, at time t, having a market share φ_{it}. Given that in the punishment phase firms make zero profit, the expected loss in this phase equals the discounted sum of expected market shares times Π

$$\delta \cdot \mathbb{E}(\varphi_{it+1} \mid \varphi_{it}) \cdot \Pi + \delta^2 \cdot \mathbb{E}(\varphi_{it+2} \mid \varphi_{it}) \cdot \Pi + \ldots = \sum_{k=1}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \cdot \Pi.$$

The gain from optimal deviation is equal to $(1 - \varphi_{it}) \cdot \Pi$. So the collusive strategy renders at least the same profit when

$$\sum_{k=1}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \geq 1 - \varphi_{it}.$$

This concludes the proof. ■

As a consequence of this result we derive analogues in our context of the results of Green and Porter (1984) and Stigler (1964). In the remainder of this section we show that, when firms experience low market shares, collusion becomes more difficult to sustain and reversion to non-collusive behavior is more likely to occur. In particular we find that the smaller a firm's market share can get, the higher the discount factor needs to be to guarantee that the trigger strategy profile is an equilibrium for all possible market share realizations. Also, since consumer loyalty reduces the volatility of market shares over time, we find in agreement with Stigler that collusion becomes easier to sustain when consumer loyalty is high.

In order to derive these analogues of Stigler and Green & Porter we need a bit more notation together with a mild assumption. Specifically, we assume for the conditional density function $f_{it+1}(\varphi_{it+1} \mid \varphi_{it})$ that $f_{it+1}(\varphi_{it+1} \mid \varphi_{it}) = 0$ outside the interval $[\underline{\varphi}, \overline{\varphi}]$ and that $f_{it+1}(\varphi_{it+1} \mid \varphi_{it}) > 0$ on the interior of the interval $[\underline{\varphi}, \overline{\varphi}]$.

Further, in accordance with the intuition that a higher market share today increases one's chances to have a higher market share in the future, we assume for the collection of cumulative probability distributions

$$F_{\varphi_{it}}(\varphi) = \int_0^{\varphi} f_{it+1}(\varphi_{it+1} \mid \varphi_{it}) \, d\varphi_{it+1}$$

that $\varphi_{it} \leq \widetilde{\varphi}_{it}$ implies $F_{\varphi_{it}}(\varphi) \geq F_{\widetilde{\varphi}_{it}}(\varphi)$ for every φ. Put slightly differently, when $\varphi_{it} \leq \widetilde{\varphi}_{it}$, the probability distribution $F_{\widetilde{\varphi}_{it}}$ of φ_{it+1} given $\widetilde{\varphi}_{it}$ stochastically dominates the probability distribution $F_{\varphi_{it}}$ of φ_{it+1} given φ_{it}. Under this assumption we can show the following fact. A sketch of its proof can be found in the Appendix.

Lemma 3.2 *For any $k \geq 1$, $\mathbb{E}(\varphi_{it+k} \mid \varphi_{it})$ is increasing in φ_{it}.*

The intuition of this result is clear. It states that the expectation of a firm's future market share is an increasing function of today's market share. In other words, higher market shares today increase the expected value of tomorrow's market share. [2] Under this condition we have the following Corollary.

Corollary 3.3 *The strategy profile T is a perfect Bayesian Nash equilibrium precisely when*

$$\sum_{k=0}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it} = \underline{\varphi}) \geq 1$$

holds for every firm i at every time t.

Proof. Due to Theorem 3.1 we know that T is a perfect Bayesian Nash equilibrium precisely when

$$\sum_{k=0}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \geq 1$$

holds for all firms i, at every time t, for any possible market share $\varphi_{it} \in [\underline{\varphi}, \overline{\varphi}]$ at time t. However, by Lemma 3.2 we know that the left-hand side of the above inequality is increasing in φ_{it}. Hence, the above inequality is satisfied for all $\varphi_{it} \in [\underline{\varphi}, \overline{\varphi}]$ precisely when it is satisfied for $\varphi_{it} = \underline{\varphi}$. ■

This is a direct analogue of the result of Green and Porter that the possibility of low market shares hamper collusion in our context. This possibility is reflected in a low value of $\underline{\varphi}$. Then, by our assumption expressed in Lemma 3.2 the left-hand side of the inequality in Corollary 3.3 is low. Thus, it will become harder to satisfy the condition for T to be a perfect Bayesian Nash equilibrium, and hence collusion is harder to sustain. It is also in agreement with Stigler (1964), since high consumer loyalty reduces the volatility of market shares. Hence high consumer loyalty increases $\underline{\varphi}$, and collusion becomes easier to sustain.

4 Collusive equilibria under public information

We now focus on the setting of Rotemberg and Saloner (1986). We show that when firms have public information on realized market shares, the incentives to deviate for a firm that has a low market share can be reduced by jointly choosing a lower collusive price. Thus, public availability of information enables firms to sustain (partial) collusion even in situations where

[2] In fact this statement itself is all we need to derive our results. However, we did not want to use this condition on expected market shares as a basic *assumption*. Instead we chose to derive it from the more basic assumption on cumulative probability distributions.

full collusion would break down. The price to pay for this enhancement of collusion is, as also argued in Rotemberg and Saloner, a lower collusive price, and hence (at least in our model with fixed aggregate demand) lower profits.

We model the phenomenon of partial collusion by slightly changing the one shot game. Since under public information the colluding firms observe the vector $\varphi_t = (\varphi_{1t}, \ldots, \varphi_{nt})$ of market shares at the start of each period t, the profit function in the one shot game can now be made contingent on the specific realization of φ_t. Before the start of the game, the colluding firms agree on a threshold level φ^* and a collusive joint profit $\Pi^* < \Pi$. The agreement is that, as long as all realized market shares φ_{it} are above the threshold φ^*, firms collude at a price level generating joint profits Π, while as soon as one or more firms have a realized market share below the threshold, firms collude at a price level that generates joint profits Π^* (partial collusion).

Effectively the collusive agreement is modeled via the one shot profit functions Π_i. Given a profile $a = (a_i)_{i \in N}$ of chosen actions, the resulting profit of firm i is denoted by $\Pi_i(a)$. The profits are now defined contingent on the realization of the vector $\varphi_t = (\varphi_{1t}, \ldots, \varphi_{nt})$.

In case $\varphi_{jt} \geq \varphi^*$ for all j. When $a_i = C$ for all i, then $\Pi_i(a) = \varphi_i \Pi$. When there is a firm k with $a_k = U$ and $a_i = C$ for all $i \neq k$, then $\Pi_k(a) = \Pi$ and $\Pi_i(a) = 0$ for all $i \neq k$. For all other strategy profiles all profits are zero.

In case $\varphi_{jt} < \varphi^*$ for some j. When $a_i = C$ for all i, then $\Pi_i(a) = \varphi_i \Pi^*$. When there is a firm k with $a_k = U$ and $a_i = C$ for all $i \neq k$, then $\Pi_k(a) = \Pi^*$ and $\Pi_i(a) = 0$ for all $i \neq k$. For all other strategy profiles all profits are zero.

Note that firms need public information on market shares in order to implement these profits.

TRIGGER STRATEGIES The trigger strategy T_i^* of firm i is defined by

$$T_i^*(h_{it}) = \begin{cases} C & \text{if all firms chose action } C \text{ in all previous rounds according to } h_{it} \\ M & \text{otherwise.} \end{cases}$$

We write $T^* = (T_i^*)_{i \in N}$ for the profile of trigger strategies. [3] Given h_t, let $p_{t+k}(h_t)$ denote the probability that $\varphi_{jt+k} \geq \varphi^*$ for all j.

Theorem 4.1 *The strategy profile T^* is a perfect Bayesian Nash equilibrium precisely when*

[3] Although the definition of T^* looks identical to the definition of the profile of trigger strategies T, due to the different information structures the set of histories on which T^* is defined differs from the set of histories on which T is defined.

for every firm i and for every information set h_t at every time t, the condition

$$\sum_{k=1}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \cdot \Big(p_{t+k}(h_t)\Pi + (1 - p_{t+k}(h_t))\Pi^*\Big) \geq (1 - \varphi_{it})\Pi$$

holds when $\varphi_{jt} \geq \varphi^*$ for all j, and the condition

$$\sum_{k=1}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \cdot \Big(p_{t+k}(h_t)\Pi + (1 - p_{t+k}(h_t))\Pi^*\Big) \geq (1 - \varphi_{it})\Pi^*$$

holds when $\varphi_{jt} < \varphi^*$ for some j.

Proof. The proof generally follows the same steps as the proofs of Theorem 3.1 and Corollary 3.3. Consider a firm i at time t with market share φ_{it}. If the firm would play U, the expected loss from the punishment period would be

$$\sum_{k=1}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \cdot \Big(p_{t+k}(h_t)\Pi + (1 - p_{t+k}(h_t))\Pi^*\Big).$$

The expected gain when $\varphi_{jt} \geq \varphi^*$ for all j is $(1 - \varphi_{it})\Pi$. So, in this case the equilibrium condition becomes

$$\sum_{k=1}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \cdot \Big(p_{t+k}(h_t)\Pi + (1 - p_{t+k}(h_t))\Pi^*\Big) \geq (1 - \varphi_{it})\Pi.$$

When $\varphi_{jt} < \varphi^*$ for some j the expected gain is $(1 - \varphi_{it})\Pi^*$. So, in this case the equilibrium condition becomes

$$\sum_{k=1}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \cdot \Big(p_{t+k}(h_t)\Pi + (1 - p_{t+k}(h_t))\Pi^*\Big) \geq (1 - \varphi_{it})\Pi^*. \qquad \blacksquare$$

The above theorem shows that in the setting with public information we can, similarly to Rotemberg and Saloner, implement partial collusion as a perfect Bayesian Nash equilibrium in trigger strategies. Also, as in R&S, strategies are based on a collusive dynamic price adjustment strategy.

This form of partial collusion under public information allows for collusion in more environments than the full collusion under private information studied in the previous section. The argument, here as well as in R&S, is straightforward: when we choose $\Pi^* = \Pi$ and $\varphi^* = \underline{\varphi}$ the above conditions are exactly equivalent to the condition in Theorem 3.1.

Thus, when full observability of market shares in not possible, this type of partial collusion cannot be implemented, and the results from Rotemberg and Saloner reduce to the results in Green and Porter. Hence, full observability enhances collusion in our model. Partial collusion incorporates the opportunities that full collusion offers, and extends to environments where full collusion is no longer sustainable.

Consumer loyalty guarantees to firms a certain fixed minimum number of consumers. Thus, when the threshold φ^* is sufficiently low, an increase in consumer loyalty tends to increase the probability $p_{t+k}(h_t)$ that $\varphi_{jt+k} \geq \varphi^*$ for all j given the history h_t. All other things being equal, this shows that an increase in consumer loyalty makes it easier to satisfy the conditions in the above Theorem, which is in line with the observations in Stigler.

5 Collusive equilibria when market shares form a martingale

An interesting special case in which we can take the above analysis a step further, and give explicit formulas for the breakdown of collusion under both private and public information, is when the stochastic process that governs the market shares forms a martingale. [4]

We start with the case of private information. In this setting, as a consequence of Theorem 3.1, we find that the trigger strategy profile is a perfect Bayesian Nash equilibrium precisely when the discount factor exceeds 1 minus the minimum market share. Thus, when the minimum market share is relatively high, and hence uncertainty is relatively low, it is easy for the firms to sustain collusion.

Corollary 5.1 *Assume private information. When the stochastic variables φ_i form a martingale, the trigger strategy profile T is a perfect Bayesian Nash equilibrium precisely when $\delta \geq 1 - \underline{\varphi}$.*

Proof. When the stochastic variables φ_i form a martingale, we have

$$\mathbb{E}(\varphi_{it+k} \mid \varphi_{it} = \underline{\varphi}) = \underline{\varphi}$$

for all t and k. [5] Thus, the equilibrium condition in Corollary 3.3 reduces to

$$\sum_{k=0}^{\infty} \delta^k \cdot \underline{\varphi} \geq 1,$$

which can be rewritten to $\delta \geq 1 - \underline{\varphi}$. ■

[4]For firm i, the stochastic process $(\varphi_{it})_{t=0}^{\infty}$ is a martingale when $\mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) = \varphi_{it}$ for every t and k. Formally, a martingale does not satisfy the condition on page 7 that $f_{it+1}(\varphi_{it+1} \mid \varphi_{it}) > 0$. However, Theorem 3.1 also holds for martingales.

[5]For a martingale it even holds that $\varphi_{it+1} = \underline{\varphi}$ with probability one when $\varphi_{it} = \underline{\varphi}$.

We now turn to the case of public information. It turns out there is an appropriate choice of φ^* such that, given δ, the adaptive trigger strategies T^* form an equilibrium whenever the profile of trigger strategies T is an equilibrium.

Corollary 5.2 *Assume public information. Suppose that the stochastic variables φ_i form a martingale. Let δ be given. Suppose further that*

$$\varphi^* \geq \frac{(1-\delta)\Pi}{\delta\Pi^* + (1-\delta)\Pi}. \tag{$*$}$$

Then the trigger strategy profile T^* is a perfect Bayesian Nash equilibrium.

Proof. Rewriting of $(*)$ yields

$$\sum_{k=1}^{\infty} \delta^k \cdot \varphi^* \cdot \Pi^* \geq (1-\varphi^*)\Pi.$$

Since the left-hand side of the inequality is increasing in φ^* and the right-hand side is decreasing, we obtain

$$\sum_{k=1}^{\infty} \delta^k \cdot \varphi_{it} \cdot \Pi^* \geq (1-\varphi_{it})\Pi$$

for all $\varphi_{it} \geq \varphi^*$. Thus, since the stochastic variables φ_i form a martingale, we find that

$$\sum_{k=1}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \cdot \Pi^* \geq (1-\varphi_{it})\Pi$$

for all $\varphi_{it} \geq \varphi^*$. Hence, since $\Pi^* < \Pi$, also

$$\sum_{k=1}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \cdot \Big(p_{t+k}(h_t)\Pi + (1-p_{t+k}(h_t))\Pi^*\Big) \geq (1-\varphi_{it})\Pi$$

for all $\varphi_{it} \geq \varphi^*$, which shows that the first set of inequalities of Theorem 4.1 is satisfied. In order to obtain the second set of inequalities, notice that the strategy profile T is a perfect Bayesian Nash equilibrium by assumption. So, by Theorem 3.1

$$\sum_{k=0}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \geq 1$$

for all t and all market shares φ_{it}. Therefore also

$$\sum_{k=1}^{\infty} \delta^k \cdot \mathbb{E}(\varphi_{it+k} \mid \varphi_{it}) \cdot \Pi^* \geq (1-\varphi_{it})\Pi^*$$

for all t and all market shares φ_{it}. The second set of conditions now follows from the observation that $\Pi^* < \Pi$. ∎

Finally note that the condition $\varphi^* \geq \frac{(1-\delta)\Pi}{\delta\Pi^* + (1-\delta)\Pi}$ can be satisfied for any given $\delta < 1$ by appropriate choices of $\varphi^* < 1$ and $\Pi^* < \Pi$.

6 Conclusion

We presented a model in which firms repeatedly engage in a Bertrand type competition model. Depending on the strategies chosen, per period profits are distributed among firms according to market shares. Market shares are allowed to fluctuate over time.

Within this model with public information on market shares we derived the conditions under which partial collusion can be implemented via trigger strategies with a dynamic price adjustment policy. Implementability of partial collusion in equilibrium is in line with the logic of Rotemberg and Saloner.

Using this model we argue that both public observability of market shares and low volatility of market shares are essential for the implementation of partial collusion. Absence of public observability prevents firms from using dynamic pricing strategies, and we revert to the logic of Green and Porter. On the other hand, in line with Stigler, we see that in our basic model with full observability, low consumer loyalty also prevents firms from using dynamic pricing strategies. Such strategies can in principle still be executed, but under low consumer loyalty, and hence high volatility of market shares, this form of partial collusion fails to satisfy the equilibrium conditions, and collusion breaks down.

Thus, our model can be seen as reconciliation of the three classical models of Stigler, Rotemberg and Saloner, and Green and Porter. We conclude that these models do not necessarily represent opposing views, but rather complement each other, and each view has its own consistent logic that indeed applies under different, mutually exclusive, conditions within our model.

7 Appendix

In this appendix we provide a sketch of the proof of Lemma 3.2. It is well known that stochastic dominance implies the following statement for monotone transformations of φ_{it+1}.

Lemma 7.1 *Let $g(\varphi_{it+1}) \geq 0$ be (strictly) increasing in φ_{it+1}. Then*

$$\mathbb{E}(g \mid \varphi_{it}) = \int g(\varphi_{it+1}) \cdot f_{it+1}(\varphi_{it+1} \mid \varphi_{it})\, d\varphi_{it+1}$$

is (strictly) increasing in φ_{it}.

Now write

$$\mathbb{E}(\varphi_{it+1} \mid \varphi_{it}) = \int \varphi_{it+1} \cdot f_{it+1}(\varphi_{it+1} \mid \varphi_{it})\, d\varphi_{it+1}.$$

By the previous Lemma, $\mathbb{E}(\varphi_{it+1} \mid \varphi_{it})$ is a strictly increasing function of φ_{it}. Thus, iterating the same argument, also

$$\mathbb{E}(\varphi_{it+2} \mid \varphi_{it}) = \mathbb{E}(\mathbb{E}(\varphi_{it+2} \mid \varphi_{it+1}) \mid \varphi_{it})$$

is a strictly increasing function of φ_{it}. In general we find that $\mathbb{E}(\varphi_{it+k} \mid \varphi_{it})$ is increasing in φ_{it} for any $k \leq 1$. This shows Lemma 3.2. ■

The above equation

$$\mathbb{E}(\varphi_{it+2} \mid \varphi_{it}) = \mathbb{E}(\mathbb{E}(\varphi_{it+2} \mid \varphi_{it+1}) \mid \varphi_{it}).$$

that we used to derive Lemma 3.2 is not a definition, it is in fact a result. In order to see why this result is true it is convenient to derive the above equation for a discrete process. Let M_1, M_2 and M_3 be three finite sets. Suppose we have transition probabilities $P(m_2 \mid m_1)$ and $P(m_3 \mid m_2)$ for all $m_1 \in M_1$, $m_2 \in M_2$, and $m_3 \in M_3$. Then

$$P(m_3 \mid m_1) = \sum_{m_2} P(m_2 \mid m_1) \cdot P(m_3 \mid m_2).$$

So,

$$\begin{aligned}
\mathbb{E}(m_3 \mid m_1) &= \sum_{m_3} P(m_3 \mid m_1) \cdot m_3 \\
&= \sum_{m_3} \sum_{m_2} P(m_2 \mid m_1) \cdot P(m_3 \mid m_2) \cdot m_3 \\
&= \sum_{m_2} \left[\sum_{m_3} m_3 \cdot P(m_3 \mid m_2) \right] \cdot P(m_2 \mid m_1) \\
&= \sum_{m_2} \mathbb{E}(m_3 \mid m_2) \cdot P(m_2 \mid m_1) \\
&= \mathbb{E}\left(\mathbb{E}(m_3 \mid m_2) \mid m_1\right).
\end{aligned}$$

The equation we used above to compute $\mathbb{E}(\varphi_{it+2} \mid \varphi_{it})$ is the continuous variant of the same result. The formula for the continuous case can be shown using the Theorem of Radon-Nikodym and Tonelli's Theorem. For further information we refer to Davidson (1994).

References

[1] Davidson J (1994) Stochastic Limit Theory. Advanced Texts in Econometrics series, Oxford University Press

[2] Ellison G (1994) Theories of cartel stability and the Joint Executive Committee. The RAND Journal of Economics 25:37–57

[3] Green EJ, Porter RH (1984) Noncooperative collusion under imperfect price competition. Econometrica 52:87–100

[4] Hendon E, Jacobsen HJ, Sloth B (1996) The one-shot deviation principle for sequential rationality. Games and Economic Behavior 12:274–282

[5] Rotemberg JJ, Saloner G (1986) A supergame-theoretic model of price wars during booms. American Economic Review 76:390–407

[6] Stigler GJ (1964) A theory of oligopoly. The Journal of Political Economy 72:44–61