Facilitating collusion by exchanging non-verifiable sales reports

David Spector

PARIS SCHOOL OF ECONOMICS
ECOLE D'ECONOMIE DE PARIS

WORKING PAPER N° 2015 – 07

Facilitating collusion by exchanging non-verifiable sales reports

David Spector

JEL Codes: C73, D83, L41

Keywords: Collusion, Communication, Politique de la concurrence, Jeux répétés

Facilitating collusion by exchanging non-verifiable sales reports

David Spector*

24 February 2015

Abstract

A number of collusive agreements involve the exchange of self-reported sales data between firms, which use them to monitor compliance with a target market share allocation. This paper shows that such communication between competitors may facilitate collusion even if all private information becomes public after a delay. The exchange of sales information may allow firms to implement incentive-compatible market share reallocation mechanisms after unexpected swings, limiting the recourse to price wars as a tool for mutual disciplining. In some cases, efficient collusion cannot occur unless firms are able to engage in such communication.

1 Introduction

The objective of this paper is to better understand the role of communication in collusive practices.

Collusion, whether tacit of explicit, requires mutual monitoring. In many recent cartel cases, monitoring took place by having companies compare each other's self-reported sales with some agreed-upon quotas, with a high frequency (often, weekly or monthly). However, these sales reports were for the most part not verifiable, at least in the short run. For instance, in several cases, reliable sales information was available only with a lag of about one year.

Prima facie, this observation is puzzling. If the goal of monitoring is to deter deviations from a collusive agreement, why couldn't a firm wanting to deviate simply undercut its competitors and misreport its sales at the same time? At

*Paris School of Economics and CNRS. Email: spector@pse.ens.fr

first glance, it seems that the only constraint on a firm's incentive to deviate is the amount of time elapsing between the date a deviation occurs and the date it is bound to being revealed to competitors, as reliable sales data become public. How, then, could the exchange of sales reports facilitate mutual monitoring and collusion if it occurs long before sales data can be verified?

Answering this question would contribute to the ongoing debate on the antitrust treatment of information exchanges. In the absence of direct evidence of cartel behavior, competition authorities face a difficult tradeoff. On the one hand, an outright ban on information exchanges would deny companies and consumers the procompetitive benefits that such exchanges may entail. Conversely, a too lenient approach would allow companies to engage in practices that could facilitate tacit collusion and harm consumers.[1]

For instance, in its guidelines on horizontal co-operation between undertakings,[2] the European Commission states that exchanges of information on past sales are not prohibited *per se* (unlike communication on future behavior) and that they should be assessed under a case-by-case approach. According to K.-U. Kühn, a former Chief Economist of the European Commission, this case-by-case approach should focus on the 'marginal impact' of the information exchanges under scrutiny on the likelihood of tacit collusion.[3]

Accordingly, this paper is an attempt to assess the marginal impact of the early disclosure of sales information long before it becomes public. When 'hard'information is publicly available in any case, should early disclosure be considered harmless cheap talk or a practice facilitating collusion relative to a no-communication benchmark?

We construct a model showing that, in a market where demand is uncertain and sales data become available with a delay, early communication on sales volumes may make collusion more efficient. Communication may reduce the recourse to price wars as a disciplining device by ensuring that unexpected market share swings are swiftly identified and compensated through short phases of market share reallocation.

Our main finding is that, for some parameter values, efficient collusion – with monopoly pricing in all periods along the equilibrium path – can occur

[1]See Kühn (2001).

[2]Official Journal of the European Union, C 11/91, 14.1.2011, Communication from the Commission — Guidelines on the applicability of Article 101 of the Treaty on the Functioning of the European Union to horizontal co-operation agreements.

[3]Kühn (2011) advocates *"an analysis of the marginal impact of the information exchange on monitoring or the scope for coordination in the market. If the marginal impact appears small, the case should be closed."*

only if communication is possible, even though such communication does not increase data verifiability.

The collusive equilibrium we derive involves no need for contact between competitors beyond the exchange of sales reports: it is symmetric, which limits the need for pre-play coordination. It is a pure-strategy equilibrium, with no need for coordination on a public randomization device. It does not involve interfirm payments.

This suggests that competition authorities should be wary of exchanges of information on past sales, even if they appear to be mere cheap talk and there is no evidence of other interfirm contacts: by itself, such communication can make tacit collusion more efficient and lead to higher prices.

The main features of the collusive equilibrium derived in this paper
The collusive equilibrium derived in this paper exhibits features similar to those observed in many recent cartels.[4]

- The collusive scheme is based on a target market share allocation. This is indeed the case in many cartels, especially those in markets in which prices are not easily observed, for instance because they are set bilaterally between sellers and buyers.

- Colluding firms exchange information on sales volumes at a high frequency. They did so every month, in the case of the lysine,[5] zinc phosphate[6] and citric acid[7] cartels, and every week in the case of the Vitamins A and E cartel[8].

- When the exchange of self-reported sales data points to a discrepancy between actual and target market shares, which can happen as a result of demand uncertainty, companies that sold above their quotas take steps to adjust their sales so as to compensate those that sold below theirs. According to Harrington (2006), whereas in some cases cartelists compensated market share swings by making payments to each other (often

[4]Harrington (2006), Levenstein and Suslow (2006).

[5]Official Journal of the European Union, L 152/24, 7.6.2001, Case COMP/36.545/F3 - Amino Acids, Decision of June 7, 2000, recital 100.

[6]Official Journal of the European Union, L 153/1, 20.6.2003, Case COMP/E- 1/37.027 - Zinc phosphate, Decision of December 11, 2001, recital 69.

[7]Official Journal of the European Union, L 239/18, 6.9.2002, Case COMP/E- 1/36.604 - Citric acid, Decision of December 5, 2001, recital 100.

[8]Official Journal of the European Union, L 6/1, 10.1.2003, Case COMP/E- 1/37.512 - Vitamins, Decision of November 21, 2001, recital 577.

under the guise of interfirm sales), in some other cases the compensation was through market share reallocations of the kind highlighted in this paper, as the continuous monitoring of sales volumes afforded colluding firms *"the opportunity to adjust their sales"*. For instance, *"the citric acid and vitamins A and E cartels engaged in 'continuous monitoring' to assess how sales matched up with quotas and, where a firm was at a pace to sell too much by the year's end, the firm was expected to slow down its sales"*. Likewise, in the zinc phosphate cartel, *"customer allocation was used as a form of compensation in the event of a company not having achieved its allocated quota."*

- Self-reported sales volumes are not instantaneously verifiable, but they can be compared to reliable data that become public with a lag. Accordingly, firms are deterred from misreporting their sales because inaccurate reports lead to price wars once they are exposed. This information structure is in line with the facts of some cartels. Reliable information can come from companies' annual reports, verification by independent auditors appointed by the cartelists, or import statistics. In many recent cases, market share information was found to become available with a delay of about one year.[9] For instance, in the lysine cartel, *"ADM reported its citric acid sales every month to a trade association, and every year, Swiss accountants audited those figures."*[10] In the Copper Plumbing Tubes cartel, import statistics were a reliable source of information on some producers' sales,[11] just like Japanese export statistics allowed companies to check Japanese manufacturers' sales reports in the sorbates cartel.[12] In the Citric Acid cartel, companies exchanged monthly information on sales, which was audited regularly, but less often than once a month.[13]

Overview of the mechanism at stake The possibility that firms could collude more efficiently by exchanging reports on their own sales may seem

[9] In our model, we make the strong assumption that the data that become publicly known after a delay are accurate and detailed, whereas in reality they may be noisy and imprecise. Firms may try to conceal some sales from auditors, as noted by Harrington and Skrzypacz (2011); and firms' annual reports, just like cross-country trade data, may not be very disaggregated.

[10] *Lysine* decision, recital 100.

[11] Official Journal of the European Union, L 192/21, 13.7.2006, Case COMP/E-1/38.069 - Copper Plumbing Tubes, Decision of 3 September 2004, recital 141.

[12] Harrington (2006), pp. 36-37.

[13] *Citric Acid* decision, recital 37.

surprising because such reports may lack credibility. Consider a collusive scheme characterized by a market share allocation and assume that, irrespective of whether communication takes place, reliable information on each firm's sales becomes public with a one-year lag. At first glance, it may seem that only this one-year lag should have an impact on the sustainability of collusion and that communication is irrelevant. If a company is requested to report its sales every month, it can still undercut its competitors and at the same time fail to report its increased sales. Both the undercutting and the lie will be detected one year later, irrespective of whether firms communicate. How could monthly reports make any difference?

The answer highlighted in this paper relies on the following difference between the deterrence of undercutting and the deterrence of misreporting: misreporting can be spotted with certainty once sales information becomes public, whereas, if demand is uncertain enough, undercutting may not be distinguishable from demand shocks.

In general, colluding firms observing market share fluctuations face a tension between two goals. On the one hand, the collusive equilibrium should deter deviations, implying that a firm gaining market share should suffer some loss thereafter. On the other hand, since market share fluctuations may be caused by demand shocks as well as by deviations, colluding firms have an interest in avoiding that such fluctuations trigger price wars.[14] One possible solution is for colluding firms to set up a mechanism that compensates for market share fluctuations while prices remain collusive, through transfers or market share reallocations. This paper considers more specifically market share reallocations: in the collusive equilibrium we derive, a firm whose market share exceeded its quota is expected to reduce its sales for a while by increasing its prices, while competitors still set the collusive price.

However, a market share reallocation phase provides weaker deterrence than a price war. Unlike a price war, it is vulnerable to deviations because the prevailing price during such a phase is collusive. A firm that is supposed to decrease its sales for a while may fail to do so and profitably undercut its competitors instead.

Therefore, in some cases, the only way to deter deviations is by having market share swings lead to price wars, even though this causes prices to fall below the monopoly level along the equilibrium path.

[14] The first formal approach to this tension is Green and Porter (1984).

In contrast, the deterrence of sales misreporting involves no such tension. A lie can be detected unambiguously by comparing hard data with earlier reports. This allows for maximal punishments following lies, in the form of price wars, without any efficiency loss since lies do not occur in equilibrium.

This contrast is the main factor making communication relevant to the effectiveness of collusion.

Absent communication, firms react to market share swings after one year, once these swings become public. Since one year of high sales at the monopoly price is very profitable, undercutting can be deterred only if it leads to painful enough consequences once market share swings become public. This is why a compensating market share reallocation may not work: in order to be dissuasive, it would have to be long, and the targeted company might have no incentive to reduce its sales for a long time. Price wars could then be the only viable deterrence mechanism.

Consider now the possibility of monthly communication on sales volumes. Since deterring lies is relatively easy - through the threat of a price war once they are exposed - firms may be induced to truthfully report their sales every month. But these monthly reports in turn make it easier for colluding firms to deter undercutting. If each firm truthfully reports its sales every month, undercutting is detected after a month rather than a year. A dissuasive market share reallocation thus need not be long. A firm having experienced an unexpectedly high market share is therefore more likely to comply with the market share reallocation mechanism. Communication could thus facilitate the recourse to market share reallocations at collusive prices, rather than to costly price wars, as a mechanism deterring deviations.[15]

The above argument holds only if there is enough uncertainty about demand. If firms cannot infer market shares from their own sales, communication may be necessary in order for swift market share reallocation to take place.

[15] In the *Citric Acid* decision, it is mentioned that the information exchange allowed colluding firms to '*monitor the correct implementation of* [*the*] *quotas and avoid, as far as possible, the need for compensation at the end of each year'* (recital 100). In that case, the outcome that colluding firms attempted to avoid thanks to the frequent exchange of sales reports was a need for compensating interfirm payments (which are ruled out by assumption in our model) rather than a price war (as per the theory presented in this paper). Despite this difference, this quote nevertheless illustrates one of the key ideas of this paper: the role of communication on sales volumes is to allow for frequent market share adjustements so as to avoid less desirable outcomes that might occur, absent communication, if market share swings were detected with a longer delay.

Relation to the literature The literature on the role of communication in collusion comprises two main sets of contributions.

Several papers address pre-play communication in the presence of private cost information.[16] The role of communication in these contexts is to facilitate the definition of an efficient collusive outcome, taking into account cost asymmetries.

Another branch of the literature addresses the role of communication as a private monitoring tool in repeated games. Compte (1998) and Kandori and Matsushima (1998) derive folk theorems, showing that infinitely patient players may achieve efficiency (*i.e.*, collusion with maximum profits) if they can communicate to facilitate mutual monitoring.[17] Our contribution differs from these on several grounds. First, the equilibria derived in these papers are highly abstract and the communication mechanisms are complex, whereas our goal is to make sense of a simple communication device, often observed in reality, namely, the exchange of own sales reports. Second, it is assumed in these papers that payments across firms are possible, whereas we rule them out. Third, these contributions do not address comparative statics, whereas we are interested in comparing the sustainability of collusion with and without communication. Fourth, the main results in these papers require the presence of at least three firms, which reflects the role of communication in their analyses: with at least three firms, each of which has some evidence on the identity of the deviator in the event of a deviation, firms can pool their information on the deviator's identity, and identify it. In contrast, our contribution focuses on firms communicating on their own sales and our main result does not require that the number of firms be at least three.

Harrington and Skrzypacz (2011, 'HS' hereafter) and Chan and Zhang (2012, 'CZ' hereafter) are closely related to this paper since they address the role of non-verifiable sales reports. They show that there may exist collusive equilibria in which companies monitor each other by exchanging sales reports, even though these reports can never be verified.

This paper differs from HS and CZ as regards both the nature of the main result and the assumptions of the model.

Whereas these two papers characterize plausible collusive equilibria involving communication, ours provides a comparative statics result, showing that for some parameter values, sustaining monopoly prices requires firms to engage in

[16] Athey and Bagwell (2001, 2008); see also Aoyagi (2002, 2007).
[17] See also Obara (2009).

communication. This result is a first step towards identifying the marginal impact on the feasibility of collusion of communication on past market outcomes.[18]

Our assumptions depart from those in HS and CZ, regarding both the information structure and the strategy space.

HS and CZ assume that neither price nor sales data ever become public, whereas we assume sales data to become public with a delay. Neither of these assumptions is inherently better than the other, since both types of information structures are found in the real world, depending on the market.

HS and CZ also assume that firms can make direct payments to each other, whereas we rule out such payments. Accordingly, in the collusive equilibria they derive, market share swings are compensated through interfirm payments, whereas in our model they are compensated through market share reallocations. Again, none of these features is inherently more relevant than the other, since both types of corrective mechanisms have been found to occur. The corrective market share reallocation mechanism highlighted in this paper, however, is probably better suited to the analysis of tacit collusion, which is the most relevant context for the analysis of information exchanges from a legal viewpoint (since in the case of explicit collusion, competition authorities can rely on evidence beyond the information exchange).

Finally, Awaya and Krishna (2014, 'AK' hereafter) prove a result related to ours, deriving conditions under which communication on past sales increases profits relative to tacit collusion. However, our model and AK's are relevant to different settings. AK assume that sales data remain private forever, and in their model, sales reports that point to a discrepancy between actual and target market shares (or to inaccurate reports) are followed by a permanent shift to non-collusive behavior. In contrast, our model is relevant to the cases described above, characterized by the following features: (i) private sales information becomes public after a delay, and (ii) sales reports that reveal a discrepancy between actual and target market shares are followed by a temporary corrective reallocation phase at collusive prices.

Organization of the paper This paper is organized as follows. After presenting the model (Section 2), we derive a sufficient condition for the nonexistence of any symmetric efficient collusive equilibrium without communica-

[18] Athey and Bagwell (2001, 2008) also derive comparative statics results on the role of communication. But they relate to a different type of communication in a different setting, namely, pre-play coordination to take into account private cost information.

tion (Section 3), and a sufficient condition for the existence of such an equilibrium when firms can exchange sales reports (Section 4). Combining these results allows us to derive a sufficient condition for such communication to be necessary to the existence of a symmetric efficient collusive equilibrium (Section 5). Section 6 concludes with a few remarks on policy implications and future research.

2 The model

2.1 Firms

There are n firms producing produce a heterogeneous good at zero cost, and facing the same rate of time preference δ. Each firm's only action, in each period, is the choice of its price (at the beginning of the period) and, in some of the games we will investigate, of its sales report (at the end of the period).

2.2 Demand

There is a continuum of consumers. Its mass is normalized to 1.

The demand function depends on the state of the world, which is drawn from some (constant) probability distribution. The draws are independent across periods. A state of the world is characterized by two parameters: total demand (which is price-inelastic), and whether demand is 'normal' or 'biased' towards one firm. Total demand can take any value $\varDelta$ within a finite set S of nonnegative numbers.[19]

There exists $V > 0$ (to be interpreted as consumers' per-unit valuation of the good) such that the demand function is as follows.

Demand in the 'normal' states of the world In a normal state of the world such that total demand is $\varDelta$, consumers consider all n goods as perfect substitutes. They attribute to the homogeneous good a subjective valuation equal to V, so that if the lowest of all firms' prices does not exceed V, total sales are equal to $\varDelta$ and they are evenly distributed among all the firms setting the lowest price. Formally, the demand function is as follows (with p_i denoting

[19] The assumption that the set of possible demand realizations is finite is made for tractability. It implies that any element of S can be equal to total demand with a strictly positive probability. Our results would carry over to a model in which the set of possible demand realizations is infinite.

the price set by Firm i and $D_i^{N,D}(p_1,..,p_n;\Delta)$ denoting Firm i's sales as a function of all firms' prices, when demand is normal and total demand is Δ):

- If $Min(p_1,...,p_n) > V$ or $p_i > Min(p_1,...,p_n)$ then $D_i^{N,D}(p_1,..,p_n;\Delta) = 0$;
- If $Min(p_1,...,p_n) \leq V$ and $p_i = Min(p_1,...,p_n)$ then $D_i^{N,D}(p_1,..,p_n;\Delta) = \frac{\Delta}{\#(\{j|p_j=Min(p_1,...,p_n)\})}$.

Demand in the 'biased' states of the world A state of the world that is biased in favor of Firm i is similar to a normal one, with one difference: consumers have a preference for Firm i's product, but a very weak one, in the sense that preferences are lexicographic. If prices are identical (and not greater than the willingness to pay), then in an i-biased demand state of the world, the entire demand goes to Firm i. But if prices are different, then the entire demand goes to the firm(s) setting the lowest price (as long as this price does not exceed the willingness to pay). This corresponds to the following demand function, with $D_j^{i,D}(p_1,..,p_n;\Delta)$ denoting Firm j's sales as a function of all firms' prices, when demand is biased in favor of Firm i and total demand is D):

- If $Min(p_1,...,p_n) > V$ or $p_j > Min(p_1,...,p_n)$ then $D_j^{i,D}(p_1,..,p_n;\Delta) = 0$;
- If $Min(p_1,...,p_n) \leq V$ and $i \neq j$ and $p_i = Min(p_1,...,p_n)$ then $D_j^{i,D}(p_1,..,p_n;\Delta) = 0$;
- If $Min(p_1,...,p_n) \leq V$ and $i \neq j$ and$p_j = Min(p_1,...,p_n)$ and $p_i \neq Min(p_1,...,p_n)$ then $D_j^{i,D}(p_1,..,p_n;\Delta) = \frac{\Delta}{\#(\{m|p_m=Min(p_1,...,p_n)\})}$;
- If $Min(p_1,...,p_n) \leq V$ and $i = j$ and $p_i = Min(p_1,...,p_n)$ then $D_j^{i,D}(p_1,..,p_n;\Delta) = \Delta$.

Notice that if demand is zero, the above characterizations of normal and biased states of the world both apply. For ease of exposition, we call a zero-demand state a normal state of the world, rather than a biased one.

Assumptions about the states of the world We make the following assumptions about the states of the world.

- We assume that for any $\Delta \in S$, the probability that the state of the world is normal with total demand equal to Δ is strictly positive.

- Let S_B ($\subset S$) denote the set of possible levels of demand in the case of biased demand (that is, the set of values of $\mathit{\Delta}$ such that, with a nonzero probability, demand is biased in favor of some firm and total demand is $\mathit{\Delta}$). For simplicity (and without any loss of generality), we assume that if $\mathit{\Delta} \in S_B$, then for any $i = \mathit{1}, ..., n$ there exists a state of the world such that demand is biased in favor of Firm i and total demand is $\mathit{\Delta}$. We also assume that the probability that demand is biased in favor of any particular firm is the same for all firms.

The two assumptions above are made for the sake of simplicity. They are not necessary for the results of this paper.

In contrast, the following assumptions, which are also made throughout the paper, are crucial because of their implications in terms of the information that firms can infer from sales data, in the event of a hypothetical collusive equilibrium characterized by all firms setting the monopoly price V.

- $0 \in S$. The possibility of a low demand shock implies that a firm selling zero cannot tell, from its own sales alone, between a low demand shock and other possible causes (such as a deviation by some competitor, or biased demand).[20]

- $S_B \neq \emptyset$. This assumption implies that in some states of the world, observing all firms' sales is not sufficient to tell whether a deviation took place. More precisely, in a state of the world (normal or biased) characterized by total demand $\mathit{\Delta} \in S_B$, then undercutting by Firm i leads to the same distribution of sales as the absence of any deviation in the presence of biased demand (in favor of Firm i).

- $S_B \supsetneq S_K$, with S_K denoting the set $\{\mathit{\Delta} | \mathit{\Delta} \in S_B and n\mathit{\Delta} \notin S\}$. This assumption implies that it is possible (whenever demand is biased and total demand belongs to $S_B \setminus S_K$) that all firms set the same price, sales are asymmetric (with only one firm having nonzero sales) and the firm making nonzero sales cannot know from observing its own sales that sales are asymmetric. The reason is that a firm benefitting from biased demand with total demand $\mathit{\Delta} \in S_B \setminus S_K$ cannot rule out, by observing its own sales, the possibility that the state of the world be normal demand with total demand equal to $n\Delta$. In other words, if $\mathit{\Delta} \in S_B \setminus S_K$, a firm that

[20] This impossibility of distinguishing between an aggregate negative demand shock and a deviation is the main element of Green and Porter's (1984) analysis.

deviated can plausibly claim that, even after observing its own sales (equal to Δ), it could not know that sales were asymmetric.

Notations We define the following probabilities.

- π^L denotes the probability that total demand is zero. The above assumptions imply that $\pi^L > 0$.
- π^H denotes the probability that total demand belongs to $S_H = S \setminus (S_B \cup \{0\})$. In the event that a firm undercuts its competitors, π^H is the probability that the information on all firms' sales is sufficient to allow one to infer that such undercutting took place. This is because, if total demand belongs to S_H, asymmetric sales can be explained only by price heterogeneity. The above assumptions imply that $\pi^H < 1$.
- π^K denotes the probability that total demand belongs to S_K. π^K can be interpreted as the probability that, after Firm i undercuts its competitors and all sales become public, Firm i will be unable to plausibly claim that it did not know that sales were asymmetric after it observed its own sales alone. The above assumptions imply that $\pi^K < 1$.
- π^B denotes the probability that demand is biased. The above assumptions imply that $\pi^B > 0$ and that the probability that demand is biased in favor of a specific set of r firms is equal to $\frac{r\pi^B}{n}$.
- D_t denotes total demand in period t (when the lowest price is lower than or equal to V).
- D denotes total expected demand ahead of any period.
- $d_t = \frac{D_t}{n}$ denotes per-firm demand in period t.
- $d = \frac{D}{n}$ denotes per-firm expected demand ahead of any period.

2.3 Timing of the game

The game is repeated for infinitely many periods, starting in period 1.

Each period is divided into three or four stages (depending on whether firms can make sales reports).

Stage 1. Firms simultaneously set prices.

Stage 2. The state of the world is determined at random.

Stage 3. Each firm observes the demand addressed to it and serves it.

Stage 4 (in some variants). Firms simultaneously make a statement on their own sales.

The above demand function and the zero cost assumption imply that the static game leads to zero prices and zero profits: the logic of Bertrand competition applies.

2.4 Information structure

The state of the world is never observed by any player.

At the end of period t, a firm observes only its own sales and (except for $t = 1$) the sales made by the other firms in period $t - 1$. In other words, at the end of a period, a firm knows all its past sales and all the sales of the other firms till the penultimate period. In the variant in which firms can communicate, a firm's information at the end of a period also includes all the statements made by all firms in all previous periods and in the current one. A firm cannot observe the prices set by the other firms either in the current period or in any past period.

As is explained above, these assumptions imply that a firm can infer little information on the state of the world, its competitors' actions, or its market share, from observing its own sales.

This difficulty in inferring market shares from own sales information is a major driver of our results. It leaves a role for communication, since pooling sales data is the only way for firms to infer market shares, even though this is in some cases not sufficient to infer the state of the world.[21] This assumption is realistic, since in many cartels the exchange of sales information was meant to allow colluding firms to estimate total sales and individual market shares, in environments characterized by the lack of price transparency.

2.5 Strategies

Strategies when communication is not possible We describe hereafter strategies in the case where communication is not possible.

Let P denote the set of possible prices (*i.e.*, the set of nonnegative real numbers).

[21] In contrast, Marshall and Marx (2008), who study a collusive equilibrium based on a market share allocation, assume that market share data are sufficient to allow a firm to identify deviations by competitors.

Let $f_t = \left(f_t^1, ..., f_t^n\right)$ denote all firms' sales at the end of period t (which become publicly observable at the end of period $t+1$).

Let p_t^i denote the price set by Firm i in period t.

For Firm i, a strategy S is an infinite set of functions $(s_1, s_2, ..., s_t, ...)$, where s_t maps all the information known to Firm X at the beginning of period t into the set of possible prices:[22]

$$s_t : \left(\left(f_1, .., f_{t-2}\right), f_{t-1}^i, \left(p_1^i, .., p_{t-1}^i\right)\right) \rightarrow p_t(\epsilon P).$$

Strategies when firms can communicate at the end of each period When communication is possible at the end of each period, strategies are modified in the following way. There is a message set M (to be specified later). At the end of each period, a firm chooses a message from M, depending on the information known to it at the end of the period (including past observed sales, past own prices, own sales in the current period, and all firms' past messages). Likewise, at the beginning of every period, the action chosen by each firm (that is, its price) depends on past own sales and prices, the other firms' publicly observed past sales, and all firms' past messages.

2.6 Equilibria

We focus throughout the paper on perfect equilibria involving pure strategies. Moreover, we restrict our attention to symmetric equilibria, as is explained hereafter. The symmetry requirement is well suited to the analysis of tacit collusion because a symmetric equilibrium is a more natural focal point than an asymmetric one.[23]

Definition - Symmetric equilibria Consider two possible series of observed sales over t periods, $f^1 = \left(\left(f_1^{11}, ..., f_1^{1n}\right), ..., \left(f_t^{11}, f_t^{1n}\right)\right)$ and $f^2 = \left(\left(f_1^{21}, ..., f_1^{2n}\right), ..., \left(f_t^{21}, ..., f_t^{2n}\right)\right)$. f^1 and f^2 are said to be symmetric with respect to σ if there exists a permutation σ of $\{1, ..., n\}$ such that for each $t' \leq t$, and each $j \in \{1, ..., n\}$, $f_{t'}^{1j} = f_{t'}^{2\sigma(j)}$. An equilibrium pair of strategies $\left(S^1, S^2\right)$ is said to be symmetric if there exists a permutation σ of $\{1, ..., n\}$ such that for any $t \geq 1$, any $j \in \{1, ..., n\}$, any pair observed sales history $\left(f^1, f^2\right)$ of length $Min(0, t-2)$ that is symmetric with

[22] We consider prices rather than probability distributions over prices because we focus on pure-strategy equilibria.

[23] Models of collusion often make this assumption (see the discussion in Athey and Bagwell (2001) and in Athey, Bagwell and Sanchirico (2004)).

respect to σ, any history of past own prices $P_t = (p_1, ..., p_{t-1})$, and any period $(t-1)$ own sales y_{t-1}:

$$s_t^{1j}\left(f^1, y_{t-1}, P_t\right) = s_t^{2\sigma(j)}\left(f^2, y_{t-1}, P_t\right).$$

Our goal is to assess whether communication can facilitate efficient collusion, that is, market outcomes such that firms earn monopoly profits. This leads us to define efficient collusive equilibria.

Definition - Efficient collusive equilibria An efficient collusive equilibrium is a subgame perfect equilibrium such that, in every period t, total profits are equal to the monopoly profit VD_t.

In the remainder of this paper, we focus on symmetric efficient collusive equilibria ("SECE").

3 The scope for collusion without communication

In this section, we derive a sufficient condition for there to exist no SECE when communication is not possible.

Proposition 1. *If the condition*

$$n - 1 - n\delta^4 - n(1-\delta)\left[\pi^H\delta^2 + \left(\pi^H\left(2 - \pi^H - \pi^K\right) + \pi^K\right)\delta^3\right] \geq 0 \qquad (1)$$

holds, then there exists no symmetric efficient collusive equilibrium.

A SECE must include a mechanism deterring deviations. This means that, after sales data revealed that only Firm i had sales above zero, outcomes must be unfavorable to Firm i for a while. The efficiency requirement mandates that, unless sales data unambiguously reveal a deviation from equilibrium behavior, such an asymmetry should be corrected not through a price war, but rather by having Firm i voluntarily withdraw from the market for some time, by setting a price strictly above the monopoly price. As we show hereafter, the constraints imposed on such correction phases by the efficiency requirement are such that a SECE strategy profile must allow any firm to undercut its competitors and earn monopoly profits in each of the first four periods (with probability 1 in the first two periods, and a smaller, but positive probability in Periods 3 and 4).

Consider a SECE, denoted Eq^*. Because of the symmetry requirement, it must be the case that in Eq^*, all firms set the monopoly price V in Period 1.

Assume that Firm 1 (without loss of generality) deviates in Period 1 and slightly undercuts its competitors. This affords Firm 1 an expected profit arbitrarily close to VD in Period 1.

At the end of Period 1, Firm 1's competitors observe that they sold zero, but they cannot infer from this information alone that a deviation took place, because zero sales could result from zero demand, or from biased demand. Therefore, in Eq^*, a firm that sold zero at the end of Period 1 must set the monopoly price V at the beginning of Period 2. Therefore, with probability 1, having slightly undercut its competitors in Period 1, Firm 1 can undercut its competitors again in Period 2, which yields an expected Period 2 profit arbitrarily close to VD.

At the beginning of Period 3, Firm 1's competitors have two additional pieces of information. They observe that they sold zero in Period 2, but this provides no information on Firm 1's Period 2 behavior, because zero sales could result from zero or biased demand. They also observed Firm 1's Period 1 sales at the end of Period 2. If $D_1 \notin S_H$, Firm 1's Period 1 sales (equal to D_1) are compatible with the possibility that Firm 1 set a price equal to V in Period 1 (if demand was zero, or biased in favor of Firm 1). Therefore, if $D_1 \notin S_H$, all the information available to Firm 1's competitors at the beginning of Period 3 is compatible with some equilibrium path. The efficiency requirement therefore implies that Firm 1's competitors set prices greater than or equal to V in Period 3. This allows Firm 1 to undercut them again and earn close to the monopoly profit. Therefore, with probability $\left(1-\pi^H\right)$, Firm 1 can earn arbitrarily close to VD in Period 3.

At the beginning of Period 4, Firm 1's competitors have two new pieces of information: their own zero sales in Period 3, and Firm 1's Period 2 sales. Their own Period 3 zero sales provide no information on Firm 1's Period 3 behavior, as zero sales could have been caused by zero demand. Consider now the information conveyed by Firm 1's Period 2 sales, in the case where $D_1 \notin S_H$, $D_1 \notin S_K$ and $D_2 \notin S_H$. Let $p(\Delta)$ denote the price which, according to Eq^*, Firm 1 should set in Period 2 after having sold Δ in Period 1. The symmetry requirement implies that $p(\Delta)$ is the price which, according to Eq^*, *any* firm that sold Δ in Period 1 should set in Period 2. If $\Delta \notin S_K$, then with a strictly positive probability demand is normal in Period 1, with total demand equal to $n\Delta$, implying that all firms sell Δ in Period 1. Eq^* thus prescribes that with a strictly positive probability, all firms set the price $p(\Delta)$ in Period 2. The efficiency requirement then implies $p(\Delta) = V$. In other words, unless $\Delta \in S_K$, Eq^* prescribes Firm 1 to set price V in Period 2 after observing that its sales were equal to Δ in Period 1. Therefore, if $D_1 \notin S_K$, Firm 1's competitors cannot infer that Firm 1 deviated simply by observing its Period 2 sales, unless $D_2 \in S_H$: Firm 1's competitors can attribute Firm 1's Period 2 sales to zero demand (if $D_2 = 0$) or to biased demand (if $D_2 \in S_B$), rather than to a deviation. If $D_1 \notin S_H$, $D_1 \notin S_K$ and $D_2 \notin S_H$, Eq^* thus prescribes that Firm 1's competitors set a price equal to V in Period 4. With probability $\left(1 - \pi^K - \pi^H\right)\left(1 - \pi^H\right)$, Firm 1 can therefore undercut its competitors in Period 4 and earn profits arbitrarily close to VD.

The above reasoning shows that if other firms set prices as prescribed by Eq^*, there exists a non-equilibrium strategy allowing Firm 1 to earn in expectation a profit arbitrarily close to VD with probability 1 in Periods 1 and 2, as well as in Period 3 with a probability of at least $\left(1 - \pi^H\right)$, and again in Period 4 with a probability of at least $\left(1 - \pi^K - \pi^H\right)\left(1 - \pi^H\right)$, leading to an expected sum of future discounted profits ahead of Period 1 arbitrarily close to (or greater than)

$$\begin{aligned} & VD\left(\tfrac{1-\delta^4}{1-\delta} - \pi^H\delta^2 - \left(1 - \left(1 - \pi^K - \pi^H\right)\left(1 - \pi^H\right)\right)\delta^3\right) \\ = \quad & nVd\left(\tfrac{1-\delta^4}{1-\delta} - \pi^H\delta^2 - \left(\pi^H\left(2 - \pi^H - \pi^K\right) + \pi^K\right)\delta^3\right) \end{aligned}$$

By the definition of an equilibrium, this should be less than or equal to the expected sum of future per-firm discounted profits ahead of Period 1 induced by compliance with Eq^*, namely $\frac{Vd}{1-\delta}$. If the former expression is greater than the latter, that is, if (1) holds, then no SECE exists.

This proof highlights the role that communication could play: each firm's ignorance of its competitors' sales prevents a compensating market share reallocation from starting right after a firm sold $\Delta \notin S_K$: a firm in such a situation cannot know whether demand was biased in its favor (which could warrant the immediate start of a correction period) or not, unless firms find a way to pool their sales data. Efficiency mandates that firms not launch a correction period till they know, but this delay increases individual incentives to deviate.

4 The scope for collusion with communication

4.1 An informal overview of the role played by communication

We consider the same environment as before, with one addition: at the end of every period, each firm is required to report its sales. Formally, each firm chooses a message at the end of each period. The meaning of the message is understood to be about each firm's sales.

The role of communication in an efficient collusive equilibrium is the following. In contrast to the situation prevailing in the absence of communication (see the above discussion of Proposition 1), communication prevents a firm from enjoying a 100% market share at the collusive price during four periods. In the presence of communication, this would be possible only for three periods.

Assume that at the end of each period, each firm is expected to report its sales, and that (i) whenever a firm is found to have lied, a price war ensues; and (ii) if reported sales point to an asymmetry, this leads to a market share reallocation phase at the expense of the firm that sold 'too much', for a few periods. Assume that Firm 1 would like to undercut its competitors in order to serve the entire demand, at the monopoly price, for as long as it can. If it does so in Period 1 and then lies about its Period 1 sales, its lie will be exposed at the end of Period 2, leading to a price war at the beginning of Period 3. In the end, Firm 1 benefits from undercutting during Periods 1 and 2 only.

Alternatively, Firm 1 could decide to truthfully report its sales. In this case, all firms would know that sales were asymmetric at the end of Period 1, and the equilibrium would prescribe a market share reallocation phase starting in Period 2, during which Firm 1 would sell zero. If Firm 1 decides to deviate and slightly undercut its competitors in Period 2, then, unless demand was zero in Period 2, its deviation is revealed at the end of Period 3, leading to a price war at the beginning of Period 4. In the end, Firm 1 benefits from undercutting in Periods 1 to 3.

In other words, in the presence of communication, a firm can deviate and obtain a 100% market share at the monopoly price for three consecutive periods at most, whereas it could do so for four periods in the absence of communication (at least for some realizations of demand). The resulting decrease in the profitability of undercutting makes efficient collusion more likely. Also, in equilibrium, the compensating market share reallocation following asymmetric sales outcomes (resulting from biased demand) starts one period after the asymmetric outcome, rather than two periods later in the absence of communication. This makes it easier to deter deviations, which facilitates collusion.

4.2 Description of the candidate equilibrium

We consider the following strategy profiles. There exists an integer k (to be interpreted as the duration of compensating market share reallocation phase after an asymmetric sales outcome is inferred from sales reports) such that, at the beginning of a period, the state of the game can be any of the following $(nk + 2)$:

- normal collusion;
- price war (which cannot occur in equilibrium);
- correction at the expense of some firm i $(1 \leq i \leq n)$, with r remaining correction periods $(1 \leq r \leq k)$.[24]

In a nutshell, 'normal collusion' is a state of the game in which all firms set the monopoly price. It gives way to temporary correction phases whenever sales reports point to asymmetric sales. During a correction phase, the firm that sold more than the other in the last normal collusion period sells zero, after which firms return to normal collusion.[25] Price wars do not occur along the equilibrium path. They occur in the (out-of-equilibrium) event of an inaccurate sales report, or of a sales report revealing that a firm failed to comply with the prescribed equilibrium strategy.[26]

Prices and messages in the candidate equilibrium Equilibrium actions at the beginning of Period t depend only on the state of the game at the beginning of Period t:

[24]In what follows, the firm at the expense of which a correction takes place in sometimes called the 'targeted firm'.

[25]As is explained below, transition rules are a bit more complex than this summary description, to account for the possibility that during a correction phase at the expense of some firm, another firm benefitted from biased demand, in which case it becomes the target of a new correction phase.

[26]This is in constrast to HS. In HS, the incentive not to misreport sales results from the existence of a function associating to each vector of sales reports the probability of a shift to a noncollusive phase. This function is such that truth-telling is an equilibrium strategy. In our model, the assumption that sales data become public with a lag allows for a simpler mechanism: lies trigger price wars once they are exposed as lies.

- If the state of the game at the beginning of Period t is 'normal collusion', then all firms set a price equal to V;
- If the state of the game at the beginning of Period t is 'collusion with a correction at the expense of Firm i', then Firm i sets a price equal to $V + 1$ whereas all other firms set a price of V;
- If the state of the game at the beginning of Period t is 'price war', then all firms set a price equal to 0;

Also, in the candidate equilibrium, each firm's message at the end of a period consists in truthfully reporting its sales, after any history along which it acted according to the candidate equilibrium in the previous periods (irrespective of the other firm's actions).[27]

Transitions between states of the game in the candidate equilibrium
The state of the game at the beginning of Period 1 is 'normal collusion'. At the beginning of Period t, the state of the game is determined as follows.

- *Evidence of lying leads to a price war.* (For $t \geq 3$). Whatever the state of the game at the beginning of Period $(t - 1)$, if the information on Period $(t - 2)$ sales, which becomes public at the end of Period $(t - 1)$, does not coincide with some firm's report (sent at the end of Period $(t - 2)$), the state of the game at the beginning of Period t is 'price war'. Also, if the sales reported by all firms at the end of Period $(t - 1)$ do not correspond to any possible outcome (for instance, if the sum of the reported sales does not belong to S), which implies that at least one firm lied, the state of the game at the beginning of Period t is 'price war'. In other words, lying, which cannot occur along an equilibrium path, leads to a price war.
- *Price wars last forever.* If the state of the game at the beginning of Period $(t - 1)$ is 'price war', then it remains 'price war' at the beginning of Period t. This reflects the fact that a price war is meant to follow outcomes that could only result from deviations. For maximal deterrence, a price war

[27] The restriction implied by the words 'after any history along which it acted according to the candidate equilibrium in the previous periods' is mentioned for the following reason: if a firm deviated from the abovementioned strategies by undercutting its competitor, it may find it optimal to also misreport its sales. For instance, on the basis of the transition rules described hereafter, if a firm undercuts its competitors in the 'normal' state of the world and demand turns out to belong to S_H, it finds it more profitable not to report its sales in order to delay the price war by one period.

lasts forever. This is compatible with efficiency since price wars cannot occur along an equilibrium path.

- *Normal collusion is followed by normal collusion if reported sales are symmetric.* If the state of the game at the beginning of Period $(t - 1)$ is 'normal collusion', none of the above conditions holds (*i.e.*, there is no evidence of lying) and all firms' reported sales for Period $(t - 1)$ are equal, then the state of the game remains 'normal collusion' at the beginning of Period t.

- *Normal collusion is followed by a correction if reported sales are asymmetric but compatible with equilibrium behavior.* If none of the above conditions holds (*i.e.*, there is no evidence of lying), the state of the game at the beginning of Period $(t - 1)$ is 'normal collusion' and for some i, Firm i's reported sales for Period $(t - 1)$ are some $\Delta \in S_B$ whereas all other firms' reported Period $(t - 1)$ sales are zero, then the state of the game at the beginning of Period t is 'correction at the expense of Firm i, with k remaining periods'.

- *Normal collusion is followed by a price war if reported sales reveal pricing that is incompatible with equilibrium behavior.* If the state of the game at the beginning of Period $(t - 1)$ is 'normal collusion' and none of the above conditions is met, then the state of the game at the beginning of Period t is 'price war'.

- *A correction phase leads to a price war if reported sales reveal a lack of compliance.* If the state of the world at the beginning of Period $(t - 1)$ is 'collusion with a correction at the expense of Firm i', and reported sales are neither such that Firm i sold zero while all others sold equal amounts adding up to some $\Delta \in S$ (which should be observed if all firms complied with the candidate equilibrium and demand was normal or biased in favor of a Firm i), nor that all firms but one sold zero, and one firm (different from Firm i) had sales belonging to S_B (which should be observed if all firms complied with the candidate equilibrium and demand was biased in favor of a firm other than Firm i), then the state of the world at the beginning of Period t is 'price war'.

- *As long as there is no evidence of lying nor lack of compliance, a correction phase continues, until it gives way to 'normal collusion'.* At the end of a correction period at the expense of Firm i with r remaining periods, if there is no evidence of lying nor lack of compliance, and the sales reported by all firms other than Firm i at the end of Period $(t-1)$ are equal, then at the beginning of period t the state of the game is 'normal collusion' if $r = 1$ and 'correction at the expense of Firm i with $(r-1)$ remaining periods' if $r > 1$. If there is no evidence of lying nor of lack of compliance, but reported sales reveal that one firm other than Firm i (labeled Firm j) had nonzero sales, whereas all others had zero sales (which can happen only if $n \geq 3$), then at the beginning of Period t the state of the world is 'correction at the expense of Firm j, with k correction periods remaining'. In other words, if a non-targeted firm benefits from biased demand, the correction at the expense of the previously targeted firm stops (even if there remained more than one period), and a new correction phase starts, targeting the firm that just benefitted from biased demand.[28]

4.3 A sufficient condition for the candidate equilibrium to be an actual equilibrium

Proposition 2. *If there are at least three firms, and there exists an integer* k *such that Conditions (2) and (3) hold, then the above strategy profiles correspond to a symmetric efficient collusive equilibrium. If there are two firms (* n=2*) and there exists an integer* k *such that Conditions (2)-(4) hold, then the above strategy profiles correspond to a symmetric efficient collusive equilibrium.*

$$n - n\delta^2 + n\pi^L\delta^2 < \delta^k \left(1 - \frac{(n-1)\pi^B}{n}\right)^k \tag{2}$$

[28]This property of the candidate equilibrium is meant for the sake of tractability. There may exist other, more complex equilibria, which make efficient collusion possible for a broader set of parameters than the candidate equilibrium outlined in this paper. However, for our purposes, it is enough to prove the existence of a SECE with communication for parameter values that are inconsistent with the existence of a SECE without communication. To this end, finding a 'good enough' SECE in the presence of communication is sufficient.

$$
\begin{aligned}
& n-1-n\delta+\delta^{k+1}\left(1-\frac{(n-1)\pi^{B}}{n}\right)^{k} \\
& -\pi^{H}\delta\left(\delta^{k}\left(1-\frac{(n-1)\pi^{B}}{n}\right)^{k}+n\delta-n\right) \\
+\pi^{L}\delta\left(1-\delta^{k}\left(1-\frac{(n-1)\pi^{B}}{n}\right)^{k}\right)+\frac{\delta^{2}\pi^{B}}{1-\delta}\left(1-\pi^{H}-\pi^{L}\right) & < 0 \qquad (3)
\end{aligned}
$$

$$1-2\delta+\delta^{k+1}+\pi^{H}\delta\left(2-\delta^{k}\right)+\pi^{L}\delta\left(1-\delta^{k}\right)>0 \qquad (4)$$

The proof of Proposition 2 is in the appendix. We explain hereafter the role of conditions (2)-(4).

Condition (2) implies that a firm expecting other firms to behave in accordance with the hypothetized equilibrium, and having set its price in accordance with the candidate equilibrium, finds that it is a best response for it to truthfully report its sales. (2) also implies that a firm subjected to a correction finds that it is optimal for it to comply with the prescribed equilibrium behavior and sell nothing for a few periods, rather than to undercut its competitors and earn monopoly profits.

Condition (3) implies that a firm not subjected to a correction has no interest in undercutting its competitors.

In the case of two firms, condition (4) implies that a firm not subjected to a correction has no interest in setting a price above its competitor's in order to sell zero and trigger a correction phase targeting its competitor.[29]

5 Comparative statics: the marginal impact of communication

Proposition 1 provides a sufficient condition for efficient collusion to be unsustainable without communication and Proposition 2 provides a sufficient condition for efficient collusion to be sustainable with communication. Combining these results leads to a sufficient condition for efficient collusion to be sustainable *only* in the presence of communication. We find that for some parameter values, communication on self-reported sales makes efficient collusion sustainaible, whereas it is not sustainable without communication - even though the underlying information structure is the same in both cases, in the sense that the timing

[29]This issue arises only in the case of two firms, as is explained in the proof.

of sales verifiability is not altered by communication:

Proposition 3. *For each* n *between 2 and 10, there exist parameter values such that no symmetric efficient collusive equilibrium exists if communication is impossible, but one such equilibrium exists if communication is possible.*

Even though statements on the set of parameters such that the existence of a SECE requires communication are difficult to interpret since the model is highly stylized, we present hereafter a few results on the set of values of δ for which collusion requires communication.

One can check that, if $n = 2$, conditions (1)-(4) are simultaneously satisfied if $\pi^H = 0.1$, π^L and π^K are close enough to zero, and $\delta = 0.82$, with $k = 2$.

Also, Graph 1 presents results for values of n between 3 and 10, in the case where π^H, π^L, π^K and π^B are close to zero. For the sake of readability, we present these results in terms of the discount rate applied over a period, that is, the rate r such that $\delta = \frac{1}{1+r}$. As is shown in Graph 1, communication substantially expands the set of discount rates compatible with the existence of a symmetric efficient collusive equilibrium (that is, discount rates such that conditions (1)-(3) are met).[30]

INSERT GRAPH 1.

6 Conclusion

The model presented in this paper casts light on the marginal impact of communication on the feasibility of collusion: the exchange of sales reports may lead to higher prices because it facilitates the recourse of colluding firms to incentive-compatible market share reallocation mechanisms, limiting the need for price wars. This effect may be present even though the information exchange is mere

[30]Since Proposition 1 states a necessary condition for the existence of a SECE without communication, and Proposition 2 states a sufficient condition for the existence of a SECE with communication, Graph 1 is conservative. The upper part of each bar depicts a set of discount rates for which communication is necessary to the existence of a SECE: it is the set of discount rates such that conditions (1)-(3) are met when π^H, π^L, π^K and π^B are close to zero. The entire set of the discount rates for which a SECE exists only when communication is possible is likely to strictly include the interval shown in Graph 1. This is why the legend for the lower part of each bar contains the word "may". For some of the corresponding values of the discount rate, collusion is possible without communication (for instance for very small discount rates, corresponding to values of δ very close to 1). But for some other values in the lower part of each bar (those close to the limit between the two parts), communication may in fact be necessary to the existence of a SECE, even though this is not visible on Graph 1.

cheap talk, in the sense that it does not make sales data verifiable sooner than they would be without communication.

While the assumption that sales data become public after a delay is strong, it nevertheless is relevant to many markets. Also, our results are robust to weakening this assumption. Assume for instance that sales data become public after being audited by third parties (as was the case in several cartels mentioned in the introduction as a motivation for this paper). If a firm that undercuts its competitors and increases its sales as a result can, with a small probability, succeed in concealing some sales from its auditor, then our results should carry over, since the profitability of undercutting is a continuous function of the probabilities of the various states of the world in the subsequent periods.

Turning to the implications for antitrust enforcement, our results suggest that the exchange of sales reports should be considered suspicious if reports revealing market share swings lead to prompt compensating movements, to an extent that cannot be explained by individual firms' unilateral profit maximization behavior, given the intertemporal pattern of demand shocks. Competition authorities should not rule out the possibility that communication is meant to facilitate such compensation mechanisms, even if the data that are exchanged between firms are not verifiable when communication takes place, and communication does not affect the date at which they will become verifiable.

Appendix

Proof of Proposition 2. Step 1: the expected sum of future discounted profits for a firm depending on the state of the game. Let $W_{c,r}$, $W_{c-,r}$ and W denote the expected sum of future discounted profits of a firm at the beginning of, respectively, a correction period at its own expense with r remaining periods ($1 \leq r \leq k$), a correction period at another firm's expense with r remaining periods, and a normal collusion period (assuming that all firms behave according to the candidate equilibrium). Since, at the beginning of a normal collusion period, the distribution of future sales is symmetric across firms, and along any equilibrium path total expected profits add up to VD, it follows that $W = \frac{VD}{n(1-\delta)}$. Also, since ahead of any period along any equilibrium path, n or $(n-1)$ firms are in a symmetric situation, $W_{c-,r} \leq \frac{VD}{(n-1)(1-\delta)}$.

Case 1: $n = 2$. In this case, it follows from the description of the transition between states of the game that $W_{c,r} = \delta^r W$.

Case 2: $n \geq 3$. In any correction period, there is a probability $\frac{(n-1)\pi^B}{n}$ that one of the non-targeted firms will benefit from biased demand, implying that in the following period, the expected sum of the previously targeted firm' future flow of discounted profits will be $W_{c-,k}$. This implies that $W_{c,1} = \delta\left[\left(1-\frac{(n-1)\pi^B}{n}\right)W + \frac{(n-1)\pi^B}{n}W_{c-,k}\right]$ and for any r between 2 and k, $W_{c,r} = \delta\left[\left(1-\frac{(n-1)\pi^B}{n}\right)W_{c,r-1} + \frac{(n-1)\pi^B}{n}W_{c-,k}\right]$. These equalities imply that for any r comprised between 1 and k,

$$W_{c,r} = \delta^r\left(1-\frac{(n-1)\pi^B}{n}\right)^r W + \frac{\frac{\delta(n-1)\pi^B}{n}}{1-\frac{\delta(n-1)\pi^B}{n}}\left[1-\delta^k\left(1-\frac{(n-1)\pi^B}{n}\right)^k\right]W_{c-,k}.$$

The above results imply that for any number of firms $n \geq 2$, the following inequalities hold:

$$\frac{\delta^r VD}{(1-\delta)n}\left(1-\frac{(n-1)\pi^B}{n}\right)^r \leq W_{c,r} \leq \frac{\delta^r VD}{(1-\delta)n}\left(1-\frac{(n-1)\pi^B}{n}\right)^r + \frac{\delta\pi^B VD}{n\,(1-\delta)}. \tag{5}$$

We prove now that for any r ($1 \leq r \leq k$), $W_{c,r} \leq W \leq W_{c-,r}$. First, the above Bellman equations characterizing each $W_{c,r}$ imply, by induction, that either for all r, $W_{c,r} \leq W$, or for all r, $W_{c,r} \geq W$. Assume (by contradiction) that $W_{c,k} > W$. Since expected profits add up to VD ahead of any period in the candidate equilibrium, the equality $W_{c,k} + (n-1)W_{c-,k} = nW$ holds, implying that $W_{c-,k} < W$. But this inequality, together with the Bellman equation characterizing $W_{c,1}$ implies that $W_{c,1} < W$, which is a contradiction.

Finally, notice that the inequality $W_{c,r} \leq W$ implies that for any r ($2 \leq r \leq k$), $W_{c,r+1} \leq W_{c,r}$ and $W_{c-,r+1} \geq W_{c-,r}$.

Step 2. Whatever the state of the world at the beginning of Period t, if Firm i complied with the strategies prescribed by the candidate equilibrium in all previous periods, and it set a price at the beginning of Period t in accordance with the candidate equilibrium, then reporting sales truthfully is a best response for Firm i, assuming that all other firms behave according to the candidate equilibrium.

Proof. Since misreporting at the end of Period t is detected at the latest at the end of Period $(t+1)$, leading to a price war from Period $(t+2)$ onwards, and expected per-period profits cannot exceed total profits VD, misreporting would allow Firm i to earn at most VD in Period $(t+1)$ and zero forever after. In contrast, complying with the strategy prescribed by the candidate equilibrium would lead, at the beginning of Period $(t+1)$, to an expected sum of future discounted profits greater than or equal to $W_{c,k}$. Since (2) implies the inequality $W_{c,r} \geq VD$ for all r, truthful sales reporting is a best response for a firm that behaved as prescribed by the candidate equilibrium.

Step 3. At the beginning of a correction period, it is a best response for all firms to set the price prescribed by the candidate equilibrium.

Proof. Assume the state of the game at the beginning of Period t is 'correction at the expense of Firm 1 (without loss of generality), with r remaining periods'. We prove first that it is optimal for Firm 1 to behave as prescribed by the candidate equilibrium, that is, by setting $p_t^1 = V+1$ and truthfully reporting its zero sales. We showed in Step 2 of this proof that conditional of setting a price equal to $V+1$ it is optimal for Firm 1 to truthfully report its zero sales. Assume that Firm 1 sets a price $p_t^1 \neq V+1$. Any price strictly greater than V yields the same zero profits as the equilibrium strategy in Period t and the same information to other firms. Therefore, such prices lead to exactly the same payoff distribution for Firm 1 in Period t and subsequent periods as $p_t^1 = V+1$. Consider now a price $p_t^1 \leq V$. Unless Period t demand is zero (which happens with probability π^L), such a price leads to nonzero sales for Firm 1, implying detection at the end of Period $(t+1)$ at the latest, and a price war starting in Period $(t+2)$ at the latest. Since Firm 1's expected per-period profit cannot exceed total profits VD, such a price leads to an expected sum of future discounted profits smaller than or equal to $VD\left[(1+\delta)+\frac{\pi^L\delta^2}{1-\delta}\right]$. (2) implies that this expression is less than $W_{c,k}$, that is, less than $W_{c,r}$ for any $r \leq k$. Therefore, it is optimal for Firm 1 to follow the strategy prescribed by the candidate equilibrium, which yields it an expected sum of future discounted profits equal to $W_{c,r}$.

Consider now a firm other than Firm 1, say, Firm 2 (without loss of generality). If there are only two firms ($n=2$) then whatever price Firm 2 sets in Period t, the state of the world at the beginning of Period $(t+1)$ is the same as the one that would prevail if Firm 2 sets a price $p_t^2 = V$. Therefore, the only effect of setting another price is to reduce Firm 2's profit. Assume now that $n \geq 3$. Complying with the actions prescribed by the candidate equilibrium leads for Firm 2 to an expected sum of future discounted profits equal to $W_{c-,r}$. Assume now that Firm 2 deviates and sets a price $p_t^2 \neq V$. If $p_t^2 > V$, then Firm 2 earns zero in Period t and whatever it reports at the end of Period t, its deviation is detected at the end of Period $(t+1)$ unless demand in Period t is zero, because a distribution of sales such that total sales are nonzero while two firms have zero sales is incompatible with equilibrium. This, together with the fact that a firm's per-period profit cannot exceed VD, implies that the corresponding expected sum of future discounted profits is less than or equal to $VD\left[\delta+\frac{\pi^L\delta^2}{1-\delta}\right]$. Condition (2) implies that this expression is less than $W_{c,k}$, and therefore less than $W_{c-,r}$. A price $p_t^2>V$ therefore cannot improve upon the behavior prescribed by the candidate equilibrium for Firm 2.

Consider now the possibility of a price $p_t^2<V$. Such a price would yield Firm 2 at most VD in Period t. If $D_t=0$, which happens with probability π^L, Period t sales are identical to what they would be absent a deviation by Firm 2. Therefore, as shown in Step 2, it would be optimal for Firm 2 in this case to report zero sales, leading in Period $(t+1)$ to either 'normal collusion' (if $r=1$) or to 'correction at the expense of Firm 1, with $(r-1)$ remaining periods' (if $r\geq 2$). This would lead at the beginning of Period $(t+1)$ to an expected sum of future discounted profits equal to $W_{c-,r-1}$ (with the slight abuse of notation $W_{c-,0}=W$). If $D_t\in S_H$, which happens with probability π^H, Firm 2's deviation is known to its competitors at the end of Period $(t+1)$ at the latest, leading to zero profits from Period $(t+2)$ onwards. Therefore, in this case, the deviation leads at the beginning of Period $(t+1)$ to an expected sum of future discounted profits smaller than or equal to VD. In all other cases ($D_t\in S_B$), a deviation leads to a sales profile that is compatible with equilibrium behavior (with demand biased in favor of Firm 2). As shown in Step 2, it is then optimal for Firm 2 to truthfully reveal its sales, leading at the beginning of Period $(t+1)$ to an expected sum of future discounted profits equal to $W_{c,k}$. Therefore, a deviation with $p_t^2<V$ would lead at the beginning of Period t to an expected sum of future discounted profits less than or equal to $VD+\delta\left(\pi^LW_{c-,r-1}+\pi^HVD+\left(1-\pi^L-\pi^H\right)W_{c,k}\right)$, whereas in the absence of deviation this expected sum is equal to $W_{c-,r}$. The difference between this expected sum in the absence and in the presence of such a deviation is thus greater than or equal to

$$\begin{aligned}
& W_{c-,r}-\left[VD+\delta\left(\pi^LW_{c-,r-1}+\pi^HVD+\left(1-\pi^L-\pi^H\right)W_{c,k}\right)\right]\\
= \quad & W-\left[VD+\delta\left(\pi^LW+\pi^HVD+\left(1-\pi^L-\pi^H\right)W_{c,k}\right)\right]\\
+ \quad & \delta\pi^L\left(W_{c-,r}-W_{c,r-1}\right)+\left(1-\delta\pi^L\right)\left(W_{c-,r}-W\right)\\
\geq \quad & W-\left[VD+\delta\left(\pi^LW+\pi^HVD+\left(1-\pi^L-\pi^H\right)W_{c,k}\right)\right].
\end{aligned}$$

Condition (3), together with inequality (5), implies that this difference is positive, so that it is a best response for Firm 2 to follow the strategy prescribed by the candidate equilibrium.

Step 4. If the state of the world at the beginning of Period t is 'normal collusion' and all firms followed the strategy prescribed by the candidate equilibrium in all previous periods, then it is a best response for Firm 1 (without loss of generality) to set $p_t^1 = V$.

Proof. Consider a subgame starting in period t, such that the state of the game at the beginning of Period t is 'normal collusion' and all firms followed the strategy prescribed by the candidate equilibrium in all previous periods. Firm 1's expected sum of future discounted profits is W if it sets price as prescribed by the candidate equilibrium (that is, $p_t^1 = V$).

We prove hereafter that, if conditions (2)-(3) hold (as well as condition (4), if $n = 2$), then in a 'normal collusion' period, neither setting a price $p_t^1 < V$ nor setting one such that $p_t^1 > V$ leads to an expected sum of future discounted profits strictly greater than W. We prove this by contradiction.

We assume that there exists a best response such that, with a positive probability, Firm 1 sets a price $p_t^1 < V$ in some 'normal collusion' state. Let W' denote Firm 1's expected sum of future discounted profits at the beginning of any such 'normal collusion' state, given this best response. Since all possible best responses must yield the same expected sum of future discounted profits at the beginning of any 'normal collusion' period, W' is also equal to Firm 1's expected sum of future discounted profits at the beginning of any subsequent 'normal collusion' period.

In Period t, setting a price $p_t^1 < V$ allows Firm 1 to serve the entire demand, leading to an expected profit below VD. If $D_t = 0$, Period t sales are identical to what they would be absent any deviation from the candidate equilibrium. Therefore, as per the proof in Step 2, it would be optimal for Firm 1 in this case to report zero sales, leading in Period $(t + 1)$ to a new 'normal collusion' period, and to an expected sum of future discounted profits equal to W' at the beginning of Period $(t + 1)$. If $D_t \in S_H$, Firm 1's deviation is known to its competitors at the end of Period $(t + 1)$ at the latest, leading to zero profits from Period $(t + 2)$ onwards. In this case, the deviation leads at the beginning of Period $(t + 1)$ to an expected sum of future discounted profits smaller than or equal to VD. In all other cases ($D_t \in S_B$), a deviation leads to sales that are compatible with equilibrium behavior (with demand biased in favor of Firm 1). The reasoning of Step 2 implies that it is then optimal for Firm 1 to truthfully reveal its sales, leading at the beginning of Period $(t + 1)$ to an expected sum of future discounted profits equal to $W_{c,k}$. Therefore, a deviation with $p_t^1 < V$ would lead at the beginning of Period t to an expected sum of future discounted profits less than or equal to $VD + \delta\left(\pi^L W' + \pi^H VD + \left(1 - \pi^L - \pi^H\right) W_{c,k}\right)$:

$$W' \leq VD + \delta\left(\pi^L W' + \pi^H VD + \left(1 - \pi^L - \pi^H\right) W_{c,k}\right),$$

implying

$$W' \leq \frac{VD\left(1 + \pi^H \delta\right) + \delta\left(1 - \pi^L - \pi^H\right) W_{c,k}}{1 - \pi^L \delta}.$$

Condition (3), together with inequality (5), implies that $W' < W$. Therefore, it is not a best response for a firm to undercut its competitors in a normal collusion period.

In order to show that there exists no best response such that, with a positive probability, $p_t^1 > V$ in some 'normal collusion' period, we need to distinguish between the case $n = 2$ and the case $n \geq 3$.

Consider first the case $n \geq 3$. If $p_t^1 > V$, then Firm 1 earns zero in Period t and whatever it reports at the end of Period t, its deviation is detected at the end of Period $(t+1)$ unless demand in Period t is zero, because a distribution of sales such that total sales are nonzero while two firms have zero sales is incompatible with equilibrium. This, together with the fact that a firm's per-period expected profit cannot exceed VD, implies that the corresponding expected sum of future discounted profits is less than or equal to $VD\left[\delta + \frac{\pi^L \delta^2}{1-\delta}\right]$. Condition (2) implies that this expression is less than W. A price $p_t^1 > V$ therefore cannot improve upon the behavior prescribed by the candidate equilibrium for Firm 1 in a 'normal collusion' period.

We consider now the case $n = 2$. We assume (by contradiction) that there exists a best response such that $p_t^1 > V$. Let Z denote Firm 1's expected sum of future discounted profits at the beginning of any 'normal collusion' period, as induced by this best response. In Period t, this price causes Firm 1's profit to be 0. If $D_t \in S_H$, Firm 1's deviation is known to Firm 2 at the end of Period t, leading to a price war and zero profits from Period $(t+1)$ onwards. In all other cases, Period t sales are compatible with equilibrium behavior (with zero demand, or demand biased in favor of Firm 2). As shown in Step 2, this implies that it is optimal for Firm 1 to truthfully report its sales. If $D_t = 0$, the state of the world at the beginning of Period $(t+1)$ is 'normal collusion'. If $D_t \in S_B$, it is 'correction at the expense of Firm 2 with k periods remaining', leading to an expected profit of $2VD$ during k periods for Firm 1, before a return to normal collusion.

Given the probabilities of the various states of the world, the expected sum of Firm 1's future discounted profits at the beginning of a 'normal collusion' period in which it sets a price $p_t^1 > V$ is thus

$$2Vd + \pi^L \delta Z + \left(1 - \pi^L - \pi^H\right) \delta^{k+1} Z + 2Vd\pi^H \delta.$$

The equality

$$Z = \pi^L \delta Z + \left(1 - \pi^L - \pi^H\right) \left(2Vd\delta \left(1 + \delta + ... + \delta^{k-1}\right) + \delta^{k+1} Z\right)$$

implies

$$Z = \frac{2\left(1 - \pi^L - \pi^H\right) Vd\delta \left(1 - \delta^k\right)}{\left(1 - \delta\right)\left(1 - \pi^L \delta - \left(1 - \pi^L - \pi^H\right)\delta^{k+1}\right)},$$

which implies in turn $Z < \frac{Vd}{1-\delta}$, as a consequence of (4). This proves that the hypothetical best response under consideration is strictly dominated by compliance with the hypothetized equilibrium. Therefore, it is not a best response for Firm 1 to set a price $p_t^1 > V$ in a 'normal collusion' period.

Step 5. The above steps imply that the strategy profile under consideration is an equilibrium strategy profile. By construction, this equilibrium is symmetric. It is also efficient, because along the equilibrium path, the states of the world that occur with a positive probability are 'normal collusion' and 'correction in favor of either firm', and in both cases, the lowest of all firms' prices is equal to the monopoly price V.

References

[1] Aoyagi, Masaki, 2002, Collusion in dynamic Bertrand oligopoly with correlated private signals and communication, Journal of Economic Theory 102, 229–248.

[2] Aoyagi, Masaki, 2007, Efficient collusion in repeated auctions with communication, Journal of Economic Theory 134, 61–92.

[3] Athey, Susan, and Kyle Bagwell, 2001, Optimal collusion with private information, The RAND Journal of Economics 32, 428–465.

[4] Athey, Susan, and Kyle Bagwell, 2008, Collusion with persistent cost shocks, Econometrica 76, 493–540.

[5] Athey, Susan, Kyle Bagwell, and Chris Sanchirico, 2004, Collusion and price rigidity, The Review of Economic Studies 71, 317–349.

[6] Awaya, Yu, and V. Krishna, 2014, On Tacit and Explicit Collusion, unpublished.

[7] Chan, Jimmy, and Wenzhang Zhang, 2012, Collusion enforcement with private information and private monitoring, Journal of Economic Theory, forthcoming.

[8] Compte, Olivier, 1998, Communication in repeated games with imperfect private monitoring, Econometrica 66, 597–626.

[9] Green, Edward J., and Robert H. Porter, 1984, Noncooperative collusion under imperfect price information, Econometrica 52, 87–100.

[10] Harrington, Joseph E., 2006, How Do Cartels Operate?, Foundations and Trends in Microeconomics 2(1), 1-108.

[11] Harrington, Joseph E., and Andrzej Skrzypacz, 2011, Private monitoring and communication in cartels: Explaining recent collusive practices, American Economic Review 101, 2425–2449.

[12] Kandori, Michihiro, and Hitoshi Matsushima, 1998, Private observation, communication and collusion, Econometrica 66, 627–652.

[13] Kühn, Kai-Uwe, 2001, Fighting collusion by regulating communication between firms, Economic Policy, 16(32), 167-204.

[14] Kühn, Kai-Uwe, 2011, Designing Competition Policy towards Information Exchanges – Looking Beyond the Possibility Results, in OECD, Information Exchanges Between Competitors Under Competition Law, DAF/COMP(2010)37.

[15] Levenstein, Margaret C., and Valerie Y. Suslow, 2006, What Determines Cartel Success?, Journal of Economic Literature, 44(1), 43–95.

[16] Marshall, Robert C., and Leslie M. Marx, 2008, Explicit Collusion and Market Share Allocation, unpublished.

[17] Obara, Ichiro, 2009, Folk theorem with communication, Journal of Economic Theory 144, 120–134.

Graph 1. The existence of a symmetric efficient collusive equilibrium according to firms' discount rate (assuming π^H, π^L, π^K, and π^B to be close to zero).