Maintaining Privacy in Cartels*

Takuo Sugaya and Alexander Wolitzky

Stanford Graduate School of Business and MIT

September 14, 2016

Abstract

It is conventional wisdom that transparency in cartels—monitoring of competitors' prices, sales, and profits—facilitates collusion. However, in several recent cases cartels have instead gone out of their way to preserve the privacy of their participants' actions and outcomes. Towards explaining this behavior, we show that cartels can sometimes sustain higher profits when actions and outcomes are only observed privately. We provide conditions under which maintaining privacy is optimal for cartels that follow the *home-market principle* of encouraging firms to act as local monopolies while refraining from competing in each other's markets. On the other hand, we also show that transparency is always optimal when firms' profit functions satisfy a certain concavity condition that holds in linear Bertrand and Cournot competition. We thus give a unified theory of both why transparency usually seems to facilitate collusion and why cartels sometimes benefit from maintaining privacy.

*For helpful comments, we thank Daron Acemoglu, Nikhil Agarwal, Yu Awaya, Kyle Bagwell, Glenn Ellison, Vijay Krishna, Kevin Li, Markus Möbius, Juan Ortner, Andy Skrzypacz, Juuso Toikka, Glen Weyl, Mike Whinston, and participants in several seminars.

1 Introduction

In the half century since the seminal paper of Stigler (1964), it has become conventional wisdom that transparency in cartels—monitoring of competitors' prices, sales, and profits—facilitates collusion. As Whinston (2006, p. 40) puts it in his monograph on antitrust economics, "Lesser observability, including more noisy signals of price cuts, makes sustaining a given supracompetitive price harder." This idea is ubiquitous in textbooks on microeconomics ("Cartel agreements are easier to enforce if detecting violations is easier"—Carlton and Perloff, 1995, p. 136) and antitrust law ("[To sustain collusion,] firms must be able to observe and compare each others' prices"—Areeda and Kaplow, 1997, p. 254). It has also been successfully applied in several well-known empirical studies, such as Albaek, Møllgaard, and Overgaard's (1997) work on the Danish ready-mixed concrete industry and Genesove and Mullin's (2001) study of the pre-war U.S. sugar industry.

There are also, however, various pieces of evidence which suggest that the conventional wisdom may not tell the whole story. Most strikingly, several recent cartels uncovered by the European Commission seem to have gone out of their way to *limit* transparency by sharing only coarse, industry-wide data, rather than the full vector of firm-level data. Harrington (2006) reports that, in the isostatic graphite cartel, this was achieved by passing around a calculator where each firm secretly entered its own sales volume, so that at the end only the sum of the reported sales was observable; firms could thus compute their own market shares, but not their competitors'. Similarly, participants in the plasterboard, copper plumbing tubes, and low density polyethylene cartels reported their individual data to a trusted intermediary (an industry group in the plasterboard case; a statistical bureau in copper plumbing tubes; a consulting firm in low density polyethylene), who then returned only aggregate statistics to the firms.[1] This behavior is a puzzle for the view that transparency facilitates collusion: as Harrington writes, "It is unclear why firms sought to maintain privacy of their market shares and to what extent effective enforcement could be achieved without market shares being commonly known among the cartel members," (p. 54).

There are also theoretical reasons why cartels might strive to maintain privacy rather than

[1] Furthermore, all of these cases concern hard-core cartels that were clearly engaged in illegal activities, so this strategy of coarsening information cannot easily be explained as an effort to comply with antitrust laws. For the details of these and other cases, see Harrington (2006) and Marshall and Marx (2012).

transparency. It is a familiar idea in economics that giving agents too much information can hurt their incentives to cooperate by giving them new ways to cheat: Hirshleifer (1971) is a classic reference. A standard example concerns a one-shot "duopoly" game, where each of two sellers must bring to a park a cart full of either ice cream or umbrellas. Ice cream is in demand on sunny days and umbrellas on rainy days, and if both sellers bring the same good they sell at a reduced price. In the absence of a weatherman, it is an equilibrium for one seller to bring ice cream, the other to bring umbrellas, and each to receive half monopoly profits in expectation. But if a weatherman tells the sellers the weather before they pack their carts, they both bring the in-demand good and split the reduced profits. Thus, in this simple example, transparency about the weather (though not transparency about the firms' actions or outcomes) actually *hinders* collusion.[2]

In this paper, we make a first attempt at investigating when the conventional wisdom that transparency facilitates collusion holds, and when on the other hand cartel participants can benefit from maintaining the privacy of their prices, sales, and profits.

The core of our analysis concerns a setting where the global market is segmented by geographic or product characteristics, and each firm has a cost advantage in its home market. In this setting, the plan of action that maximizes cartel profits has each firm price optimally in its home market and refrain from entering its competitors' home markets. This situation is typical in many industries (though it is perhaps less well-studied in theoretical models of collusion). For example, in the choline chloride cartel, which consisted of firms based in Europe (Akzo Nobel from the Netherlands; BASF from Germany; UCB from Belgium) and North America (Bio Products and DuCoa from the US; Chinook from Canada), the agreement was that the European firms would exit the North American market while the North American firms would exit Europe. Harrington refers to this plan as the "home-market principle," and writes, "A common principle to a number of cartels was the 'home-market principle,' whereby cartel members would reduce supply in each other's home markets," (Harrington, 2006, p. 34).[3]

[2]This story was told to us by Faruk Gul. Gul attributes it to Howard Raiffa, who in turn apparently attributed it to Hirshleifer.

[3]Harrington documents the home-market principle in a wide range of industrial cartels, including the isostatic graphite and copper plumbing tubes cartels discussed above. See Section 3 for further discussion. Another example is perhaps given by the Texas school milk cartel studied by Pesendorfer (2000). Pesendorfer resports that this cartel operated by dividing the market based on cost advantages and refraining from entry

Our first result provides conditions under which, when demand and cost conditions change over time (so the environment is given by a *stochastic* game), the home-market principle is easier to sustain when firms observe their competitors' prices and sales *less* precisely (in the sense of Blackwell, 1951). The idea is as follows: When the cartel divides the market, each firm may not need to know the demand state in the "foreign" markets to price optimally in its home market. However, revealing informative past behavior in the foreign markets (which amounts to revealing information about the demand state in the foreign markets) does help the firm tailor potential deviations—in which it violates the collusive agreement by entering the foreign markets—to the current conditions in these markets. Sustaining the home-market principle therefore requires more patience when firms observe their competitors' actions more precisely. The conventional wisdom that transparency facilitates collusion thus fails badly in stochastic environments where the cartel tries to segment the market.

We also provide some suggestive narrative evidence that our results can help explain some cartels' efforts to maintain the privacy of their members' firm-level data. In particular, we discuss several cases drawn from the European Commission decisions analyzed by Harrington (2006) and Marshall and Marx (2012), and show that a number of cartels relied on the combination of the home-market principle and coarse information-exchange suggested by our model.

Finally, we present two additional sets of theoretical results aimed at clarifying the boundaries of our core model and mechanism. We first investigate whether maintaining privacy can ever be required for sustaining collusion in a standard *repeated* game setting, with no payoff-relevant state variables. In this canonical setting, it is much less clear whether transparency can ever hinder collusion: for example, Kandori (1992) has shown that, when one restricts attention to *public* equilibria of repeated games with imperfect *public* monitoring, improved observability can only expand the equilibrium payoff set, consistent with the conventional wisdom. Nonetheless, we show that maintaining privacy can again be essential for supporting collusion, even in the simplest entry deterrence game.

Specifically, we again consider the home-market principle, but now focus on the particular market structure where one home firm (acting as an incumbent according to the home-market principle) faces the potential entry of a foreign firm. Costs and demand are

in others' home areas, while occasionally swapping contracts among firms in response to changing costs.

stationary, so there is no exogenous source of time variation. We model this situation as a canonical repeated entry-deterrence game, and show that in this game firms cannot deter inefficient entry into their home markets when they perfectly observe each other's actions, but that entry can be deterred if the firms observe only their own actions and profits, while a mediator (like the consultant or industry group in the examples above) collects this information from the firms and returns only the coarsest aggregate information required to detect deviations. The intuition is that—despite the stationarity of the physical environment—a collusive equilibrium must sometimes have a non-stationary path of play, where some histories represent more tempting times for a firm to deviate than others. In these cases, revealing too many details of the history can prompt deviations.

Second, we ask whether there are any broad classes of markets in which maintaining privacy *cannot* help sustain collusion. Our last set of results identifies one such class: for markets in which, (1) firms' profits are jointly *concave* in their own and their competitors' actions, (2) each firm's maximum instantaneous deviation gain is jointly *convex* in actions, and (3) a firm that deviates can be held to its minmax payoff, we show that no information structure can sustain higher industry profits than perfect monitoring (i.e., full transparency). For example, these conditions are satisfied by Cournot or differentiated-product Bertrand competition with linear demand curves. This result follows because—as we show—these conditions are sufficient for the best collusive equilibrium to have a deterministic path of play (roughly, if profits are concave and deviation gains are convex, stochastic arrangements decrease the gain from collusion and increase the temptation to deviate), and in a deterministic equilibrium transparency does not give potential deviators any actionable information.

Putting all of our results together, we find that transparency about actions and/or outcomes can hinder collusion, both in stochastic environments where firms follow the home-market principle and in repeated entry-deterrence games, but that on the other hand the conventional wisdom that transparency helps collusion is borne out in repeated games with sufficiently well-behaved stage games. This gives a possible theory both for why the conventional wisdom that transparency facilitates collusion has proved to be so important and for why it occasionally seems to be contradicted by cartel practices.

Finally, we note a broader implication of our analysis for the problem of managing long-run relationships. In a "smooth" relationship—one with a deterministic, "cooperative" equi-

librium path of play—perfect monitoring of actions is the best information structure, because in a smooth relationship all the information structure needs to do is detect deviations. So, if a smooth relationship is optimal (which we show is the case for concave games), then perfect monitoring is optimal. However, in a "rocky" relationship—one where actions, and in particular the level of cooperation, fluctuate randomly—perfect monitoring may be suboptimal, as there is now a benefit from keeping players in the dark regarding the current state of the relationship. Thus, if a rocky relationship is optimal, then maintaining privacy can be valuable.[4]

The reader may wonder how, fifty years after Stigler's paper, we can describe our paper as a "first attempt" at investigating whether transparency facilitates collusion. The answer is that the overwhelming majority of the literature on collusion either assumes that monitoring of actions is perfect (Friedman, 1971; Abreu, 1986; Rotemberg and Saloner, 1986) or assumes that monitoring is imperfect but restricts attention to equilibria where firms condition their actions only on publicly available information (Green and Porter, 1984; Abreu, Pearce, and Stacchetti, 1986; Athey and Bagwell, 2001). In the latter case, as mentioned above, Kandori (1992) shows that improved observability can only help collusion: the intuition is that, as signals become more precise, the firms always have the option of simply agreeing to condition their play on a "noised up" version of the signals.[5] However, in the more realistic case where firms receive private signals, this intuition break down completely, as there is no way to force a firm to condition only on a noised up version of its private information. Thus, to understand whether improved observability helps or hinders collusion, one must consider repeated games with private monitoring, such as Stigler's original secret price-cutting game.

Among the relatively few papers that have studied collusion with private monitoring, several focus on the "folk theorem" question of providing conditions for first-best collusion

[4] A recent paper by Bernheim and Madsen (2016) argues that relationships that are in this sense rocky are the norm in industrial cartels. Among the cartels in which they report significant discord are those in lysine, nucleotides, and citric acid. Restricting attention to stationary equilibria, Bernheim and Madsen provide a complementary model of rocky relationships in cartels, where monitoring is perfect but the presence of fixed costs of operating in rivals' home markets leads to an equilibrium in mixed strategies.

[5] The result of Abreu, Milgrom, and Pearce (1991) that delaying the arrival of information can reduce the scope for deviations and increase efficiency involves the consideration of private strategies in a repeated game with imperfect public monitoring. Like our results, this finding is based on the idea that pooling information sets can be good for incentives. But there are also many differences between the results. For example, their result restricts attention to strongly symmetric equilibria, and their model and result are unrelated to the home-market principle.

to be sustainable when the firms are sufficiently patient (Aoyagi, 2007; Hörner and Jamison, 2007). Another set of papers asks when letting firms communicate is necessary or sufficient for sustaining collusion with private monitoring (Athey and Bagwell, 2001; Aoyagi, 2002; Harrington and Skrzypacz, 2011; Rahman, 2014; Awaya and Krishna, 2015). In particular, Rahman and Awaya and Krishna show that communication may be necessary for sustaining collusion if the quality of monitoring is sufficiently poor. This occurs because, in their models, communication can essentially be used to improve the precision of monitoring, as in the earlier papers of Compte (1998) and Kandori and Matsushima (1998). However, none of these papers address our question of whether monitoring can be *too* precise, in that worse observability can actually help sustain collusion.[6,7]

A companion paper (Sugaya and Wolitzky, 2016) contains an example where players in a repeated game benefit from imperfections in the monitoring technology.[8] The example in that paper is somewhat similar to the entry deterrence example in the current paper; however, it is substantially more complicated and does not have a clear connection to collusion or entry deterrence.[9] Moreover, the main result of the companion paper is instead a sufficient condition for transparency (i.e., perfect monitoring) to be optimal, like Proposition 7 of the current paper. The result is however completely different, as it applies to all two-player games for sufficiently high discount factors with a mediator, rather than to "concave" n-player games at any discount factor with or without a mediator.

The remainder of the paper is organized as follows. Section 2 presents our main result: transparency hinders collusion in stochastic games that satisfy the home-market principle. Section 3 discusses several real-world cartels through the lens of this model. Sections 4 and 5 contain our additional theoretical results: Section 4 shows that transparency also hinders

[6] Athey and Bagwell (2001) do present a numerical example where firms benefit from limiting communication about payoff-relevant state variables.

[7] Another feature of Harrington and Skrzypacz (2011) is that they focus on equilibria that are essentially stationary (and hence deterministic), with incentives provided by transfers among firms. Goldlücke and Kranz (2012) show that, within the class of perfect public equilibria, restricting attention to stationary equilibria is without loss of generality when transfer are available. This result would not apply to our model even if we allowed for transfers (we do not), as we do not restrict attention to public equilibria.

[8] See Kandori (1991a), Sekiguchi (2002), Mailath, Matthews and Sekiguchi (2002), and Miyahara and Sekiguchi (2013) for examples where players benefit from imperfections in monitoring in finitely repeated games, due to a somewhat different mechanism.

[9] On the other hand, the example in the companion paper does not require the players to have access to a mediator.

collusion in the repeated entry-deterrence game; and Section 5 shows that transparency always helps support collusion if the underlying stage game satisfies the concavity/convexity conditions described above. Section 6 concludes.

2 Transparency can Hinder Collusion in Stochastic Games: The Home-Market Principle

We begin with our main explanation for the efforts of some cartels to preserve the privacy of their participants' actions: under the home-market principle, collusion is easier to sustain when firms have *less* precise information about their competitors' prices, costs, or sales. We also provide a parameterized example in which the optimal degree of transparency is increasing in the discount factor, so that maintaining privacy is most valuable when the firms are just patient enough to sustain collusion.

2.1 Model

There are n firms competing in n distinct markets. The markets can represent niches in geographic or product attribute space, or can correspond to n large consumers who comprise the demand side of the market. In every period $t = 0, 1, 2, \ldots$, each firm i can produce in market j at constant marginal cost $c_i^j \geq 0$, where the vector of *cost states* $c = \left(c_i^j\right) \in \left(C_i^j\right) = C$ can change over time as described below. We will assume that $c_i^i \leq c_j^i$ for all $i \neq j$ and all $c \in C$: that is, firm i has a cost advantage in its corresponding "home" market, market i.[10]

In every period t, firm i chooses a price vector $\left(p_i^j\right)_{j=1}^n \in \left(\mathbb{R}_+ \cup \{\infty\}\right)^n$, where p_i^j is firm i's price in market j. As we will see, setting $p_i^j = \infty$ corresponds to "staying out" of market j, or equivalently setting a price so high that no consumer will ever purchase. Let $\underline{p}^j = \min_i p_i^j$ be the lowest price in market j. Demand in market j is given by a function $D\left(\underline{p}^j, s^j\right)$, where $s^j \in S^j$ is the current *demand state* in market j and S^j is the set of possible market j demand states. Let $s \in S = \left(S^j\right)_{j=1}^n$ denote a vector of market demand states. Assume that the function $D\left(p^j, s^j\right)$ is continuous, non-negative, and strictly increasing in s^j,

[10] For a recent model of dynamic price competition with a similar distinction between home and foreign markets, see Bernheim and Madsen (2016).

with $D\left(p^j, s^j\right) p^j$ bounded and $D\left(\infty, s^j\right) = 0$.[11] The lowest-price firms in market j supply all $D\left(\underline{p}^j, s^j\right)$ units at price $\underline{p}^j$, with the market allocated to the home firm in case of a tie (or allocated equally among the lowest-price firms if the home firm does not have the lowest price).[12] Denote the vector of sales at price vector p and demand state vector s by $q\left(p, s\right) = \left(q_i^j\left(p^j, s^j\right)\right)$. Finally, assume that the vector $\left(c_{t+1}, s_t\right) \in C \times S$ of period $t+1$ cost and period t demand states follows a Markov process, with Markov transition function $M : C \times S \to \Delta\left(C \times S\right)$.

At the end of period t, firms observe their own period t prices, period t sales, and period $t+1$ costs (next period costs), and can also receive signals about the entire vector of period t prices, period t demand states, and period $t+1$ costs (see below about the initial period). Specifically, there is a finite set of signals $Z = \left(Z_i\right)$ and a family of conditional probability distributions (or "information structure") on Z, $\pi\left(z|p, c', s\right)$, such that signal z is realized with probability $\pi\left(z|p, c', s\right)$ when p is the vector of prices and $\left(c', s\right)$ is the vector of next period cost and current demand states. When signal z is realized, firm i observes only its i^{th} component, z_i. Letting π_i denote the marginal distribution over z_i, we impose the following key assumption, which says that a firm can distinguish between situations in which another firm enters its home market and situations in which this does not occur:

Assumption 1 Fix an arbitrary market j and price vectors p and $\hat{p}$ such that $p_i^j = \infty$ for all $i \neq j$ and $\hat{p}_i^j \neq \infty$ for some $i \neq j$. Then, for every $z_j \in Z_j$, if $\pi_j\left(z_j|p, c', s\right) > 0$ for some $\left(c', s\right) \in C \times S$ then $\pi_j\left(z_j|\hat{p}, c', s\right) = 0$ for all $\left(c', s\right) \in C \times S$; and if $\pi_j\left(z_j|\hat{p}, c', s\right) > 0$ for some $\left(c', s\right) \in C \times S$ then $\pi_j\left(z_j|p, c', s\right) = 0$ for all $\left(c', s\right) \in C \times S$.

In effect, Assumption 1 says that firm i can behave so passively in market j that with probability 1 firm j cannot misinterpret its behavior as an attempt to steal the market. Another interpretation is that firm i can certify to firm j that it has not tried to sell in market j. As we discuss further below, we believe that this assumption is reasonably consistent with applications of the home-market principle in practice, where market segmentation by geography or by large consumers greatly alleviates the difficulty of monitoring entry in one's home market. For example, Harrington writes that "An attraction to a customer allocation

[11] We follow the convention $0 \cdot \infty = 0$.

[12] This tie-breaking rule would emerge endogenously in Bertrand competition with heterogeneous costs if consumers could choose from whom to purchase.

scheme is that monitoring is relatively easy since, if a firm was to supply a particular buyer, it would surely know whether that buyer ended up buying from someone else," (2006, p. 46).

We also assume that signals have a product structure, which allows every signal z_i to be decomposed into signals $x^k_{j,i}$ of individual prices ("i's signal of j's price in market k") and a signal y_i of the joint vector of current sales and next period cost. This means that signals of some prices are not directly informative about either other prices or the demand and cost states—though of course in equilibrium firms will draw inferences about demand and costs based on price signals.

Assumption 2 There exist finite sets $X^k_j = \left(X^k_{j,i}\right)$ and $Y = (Y_i)$ and families of conditional probability distributions $\pi^{X^k_j}\left(x^k_j | p^k_j\right)$ and $\pi^Y\left(y | p, c', s\right)$ such that

1. $Z = \left(\prod_{j,k} X\right) \times Y$,
2. $\pi\left((x,y)\,|p,c',s\right) = \left(\prod_{j,k} \pi^{X^k_j}\left(x^k_j | p^k_j\right)\right) \pi^Y\left(y|p,c',s\right)$, and
3. $\pi^Y\left(y|p,c',s\right) = \pi^Y\left(y|\hat{p},c',s\right)$ whenever $q\left(p,s\right) = q\left(\hat{p},s\right)$.

Finally, at the beginning of period 0, each firm i knows its own cost $c_{i,0}$. We assume that the distribution of the current state profile $\left(s_{i,0}\right)_i$ is determined by $\left(c_{i,0}\right)_i$. (For simplicity, we assume that firms do not observe additional signals in the initial period, except for Assumption 4 below.)

The firms have a common discount factor $\delta < 1$. The solution concept is sequential equilibrium.

2.2 Conditions under which Transparency Hinders Collusion

Our first result provides conditions under which transparency weakly hinders collusion, in that firms must be more patient in order to sustain the first-best collusive scheme under a more informative information structure. In addition to the maintained Assumptions 1 and 2, the first condition is that demand transitions at the level of an individual firm or market depend only on that firm/market's current demand.[13] The economic content of this

[13] Note that this assumption does not imply that state transitions are independent across firms or markets.

assumption is that, if a firm knows both its own costs and the demand state in its home market, it does not require information about costs and demand in other markets in order to price optimally in its own market.

Assumption 3 For every market j there is a function $M_j^S : S^j \to \Delta(S^j)$ such that $M_j^S(s^j) = M(c', s)|_{S^j}$ for all $(c', s) \in C \times S$, where $M(c', s)|_{S^j}$ denotes the projection of $M(c', s)$ onto S^j.

[*we can allow s^j to depend on c'_j since now c'_j is observable for j; it relaxes the assumption, but only a bit, I think*]

If firm i were a monopoly in market i and observed the period t vector of demand states s_t and next period cost vector c_{t+1}, it would set price p_i^i in period $t+1$ to maximize $\mathbb{E}\left[D\left(p_i^i, s_{t+1}^i\right)\left(p_i^i - c_{i,t+1}^i\right)|c_{t+1}, s_t\right]$, which equals $\mathbb{E}\left[D\left(p_i^i, s_{t+1}^i\right)\left(p_i^i - c_{i,t+1}^i\right)|c_{i,t+1}^i, s_t^i\right]$ by Assumption 3. Let $p_i^m\left(c_{i,t+1}^i, s_t^i\right)$ be a solution to this problem. Let $p_i^m(c_{i,}, \emptyset)$ be a maximizer of $\mathbb{E}\left[D\left(p_i^i, s_0^i\right)\left(p_i^i - c_{i,0}^i\right)|c_{i,0}\right]$.

Due to the cost advantage of producing in one's home market, there is essentially a unique joint plan of action that sustains first-best expected industry profits. Borrowing Harrington's terminology, we refer to this action plan as the *home-market principle*:

- Firms only price competitively in their home markets: in every period, $p_j^i \geq p_i^m\left(c_{i,t}^i, s_{t-1}^i\right)$ for all $\left(c_{i,t}^i, s_{t-1}^i\right) \in C_i^i \times S^i$. Given this, each firm can perfectly infer the previous demand state in its home market in every period $t > 0$ (by the assumption that $D\left(\underline{p}^i, s_{t-1}^i\right)$ is strictly increasing in s_{t-1}^i).
- In period 0, each firm i sets price $p_i^m\left(c_{i,1}^i, \emptyset\right)$ in its home market.
- In period $t > 0$, each firm i sets price $p_i^m\left(c_{i,t}^i, s_{t-1}^i\right)$ in its home market.

The second and final condition is that costs are observable. This condition lets firms punish a deviator as harshly as possible by pricing at its cost in its home. (We also present a version of our result without this assumption below.)

Assumption 4 $\pi_i(z|p, c', s)\,\pi_i(z|p, \hat{c}', s) = 0$ for all $c' \neq \hat{c}'$ and all (i, z, p, s). Moreover, the initial period cost profile c_0 is observable to each firm.

Recall that an information structure π' is *(Blackwell) more informative* than π for every firm i if for all i there exists a function $f_i : Z_i \times Z_i \to [0,1]$ such that $\sum_{z_i \in Z_i} f_i(z_i, z_i') = 1$ for all $z_i' \in Z_i$ and, for all $z_i \in Z_i$, p, c', and s,

$$\pi_i(z_i|p, c', s) = \sum_{z_i' \in Z_i} f_i(z_i, z_i')\, \pi_i'(z_i'|p, c', s).$$

Proposition 1 *Under Assumptions 1–4, for any information structure π there is a cutoff discount factor $\delta^*(\pi)$ such that the home-market principle (and thus first-best industry profits) can be sustained in equilibrium if and only if $\delta \geq \delta^*(\pi)$; furthermore, if π' is more information than π then $\delta^*(\pi') \geq \delta^*(\pi)$. In this sense, a more informative information structure hinders collusion.*

Proof. We show that if there is an equilibrium σ' satisfying the home-market principle with information structure π', then with any less informative information structure π there is an equilibrium where firms follow the home-market principle while always pricing at ∞ outside their home markets (on path). To see this, consider the strategy profile σ where firms follow the home-market principle while always pricing at ∞ outside their home markets, and if any firm i either enters a foreign market or detects entry in its home market it sets price $p_{i,t}^j = c_{j,t}^j$ for all j in every subsequent period t. (Note that this feasible by Assumption 4.) This strategy profile holds a deviator to its minmax payoff of 0. It is also sequentially rational at off-path histories, because firms are willing to price below cost outside their home markets as they make no sales at these prices due to the tie-breaking rule. It thus remains only to show that this strategy profile is sequentially rational at on-path histories.

By Assumption 3, firm i's on-path period t continuation payoff (i.e., its expectation at the beginning of period t of its payoff starting in period $t+1$) depends only on its current cost $c_{i,t}^i$ and the previous demand state in its home market s_{t-1}^i, which firm i can perfectly infer in any equilibrium satisfying the home market principle, for any information structure. Denote this continuation payoff by $V_i\left(c_{i,t}^i, s_{t-1}^i\right)$. Next, denote firm i's maximum gain from a deviation at history $h_i^t = (p_{i,\tau}, c_{i,0}, c_{i,\tau+1}, q_{i,\tau}, x_{i,\tau}, y_{i,\tau})_{\tau=0}^{t-1}$ under strategy profile σ' and

information structure π' by

$$d_i^{\sigma',\pi'}\left(h_i^t\right)=\sum_{j\neq i}\max_{p_i^j}\Pr^{\sigma',\pi'}\left(p_i^j<p_{j,t}^j|h_i^t\right)\mathbb{E}^{\sigma',\pi'}\left[D\left(p_i^j,s_t^j\right)\left(p_i^j-c_{i,t}^j\right)|h_i^t,p_i^j<p_{j,t}^j\right].$$

As σ' is an equilibrium with information structure π' and each firm's minmax payoff is 0, we have $\delta V_i\left(c_{i,t}^i,s_{t-1}^i\right)\geq(1-\delta)\,d_i^{\sigma',\pi'}\left(h_i^t\right)$ for all i, $c_{i,t}^i$, and s_{t-1}^i, and all histories h_i^t that lie on path under strategy profile σ' with information structure π' and are consistent with $\left(c_{i,t}^i,s_{t-1}^i\right)$. Denote this set of histories by $H_i^t\left(\sigma',\pi',c_{i,t}^i,s_{t-1}^i\right)$. On the other hand, the strategy profile described above is sequentially rational if $\delta V_i\left(c_{i,t}^i,s_{t-1}^i\right)\geq(1-\delta)\,d_i^{\sigma,\pi}\left(h_i^t\right)$ for all i, $c_{i,t}^i$, and s_{t-1}^i, and all histories $h_i^t\in H_i^t\left(\sigma,\pi,c_{i,t}^i,s_{t-1}^i\right)$, the set of on-path histories with strategy profile σ and information structure π consistent with $\left(c_{i,t}^i,s_{t-1}^i\right)$. Hence, it suffices to show that $\sup_{h_i^t\in H_i^t\left(\sigma',\pi',c_{i,t}^i,s_{t-1}^i\right)}d_i^{\sigma',\pi'}\left(h_i^t\right)\geq\sup_{h_i^t\in H_i^t\left(\sigma,\pi,c_{i,t}^i,s_{t-1}^i\right)}d_i^{\sigma,\pi}\left(h_i^t\right)$ for all firms i and all $\left(c_{i,t}^i,s_{t-1}^i\right)\in C_i^i\times S_i$.

Let $\chi_i^t\left(\sigma,\pi\right)$ be the distribution over histories h_i^t under strategy profile σ with information structure π. Let $\chi_i^t|_{\left(X_j^j\right)\times Y}\left(\sigma,\pi\right)$ be projection of $\chi_i^t\left(\sigma,\pi\right)$ onto $\left(\left(X_{j,i}^j\right)_{j=1}^n\times Y_i^\tau\right)_{\tau=0}^{t-1}$. For any strategy profile $\hat{\sigma}$ satisfying the home-market principle and any information structure $\hat{\pi}$, $d_i^{\hat{\sigma},\hat{\pi}}\left(h_i^t\right)$ depends on h_i^t only through firm i's belief about the vector of cost $\left(c_t^j\right)_j$ and demand state $\left(s_{t-1}^j\right)_j$, as the distribution over both firm i's opponents' period t actions $\left(p_{j,t}\right)_j$ and the period t vector of demand states $\left(s_t^j\right)_j$ are determined by $\left(c_{j,t},s_{t-1}^j\right)_j$. Under strategy profile σ, as all firms price at ∞ outside their home markets, at any on-path history firm i's beliefs about $\left(c_t^j,s_{t-1}^j\right)_j$ depend only on $\left(c_t^j\right)_j$, s_{t-1}^i, $\left(\left(x_{j,i}^j\right)_{j=1}^n\right)_{\tau=0}^{t-1}$, and $\left(y_i^\tau\right)_{\tau=0}^{t-1}$ (by Assumption 2).[14] Now, the assumption that π' is more informative than π for i implies that $\chi_i^t|_{\left(X_j^j\right)\times Y}\left(\sigma,\pi'\right)$ is more informative than $\chi_i^t|_{\left(X_j^j\right)\times Y}\left(\sigma,\pi\right)$ (viewing $\chi_i^t|_{\left(X_j^j\right)\times Y}\left(\sigma,\pi'\right)$ and $\chi_i^t|_{\left(X_j^j\right)\times Y}\left(\sigma,\pi\right)$ as mappings from realizations of $\left(c_0,\left(\left(c_{j,\tau+1},s_\tau^j\right)_j\right)_{\tau=0}^{t-1}\right)$ to distributions over $\left(\left(X_{j,i}^j\right)_{j=1}^n\times Y_i^\tau\right)_{\tau=0}^{t-1}$); this may be proved by induction on t. It is immediate from Bayes' rule that if $\chi_i^t|_{\left(X_j^j\right)\times Y}\left(\sigma,\pi'\right)$ is more informative than $\chi_i^t|_{\left(X_j^j\right)\times Y}\left(\sigma,\pi\right)$, then firm i's belief about any $\left(\left(X_{j,i}^j\right)_{j=1}^n\times Y_i^\tau\right)_{\tau=0}^{t-1}$-measurable random variable after any observation from $\chi_i^t|_{\left(X_j^j\right)\times Y}\left(\sigma,\pi\right)$ is a convex combination of its beliefs after various observations from

[14] Under Assumption 4, a firm's belief about $\left(c_{j,t}\right)_j$ is degenerate, so we could omit this argument. We include it to highlight that the same argument applies for Proposition 2, where Assumption 4 is dropped.

$\chi_i^t|_{\left(X_j^j\right)\times Y}(\sigma,\pi')$. Hence, firm i's belief about $\left(c_{j,t},s_{t-1}^j\right)_j$ at any history $h_i^t\in H_i^t(\sigma,\pi)$ under strategy profile σ with information structure π is a convex combination of its beliefs about $\left(c_{j,t},s_{t-1}^j\right)_j$ at various histories $h_i^t\in H_i^t(\sigma,\pi')$ under strategy profile σ with information structure π'. In addition, for every realization of $\left(c_0,\left(\left(c_{j,\tau+1},s_\tau^j\right)_j\right)_{\tau=0}^{t-1}\right)$, realized home-market prices and sales are the same under σ and σ', so $\chi_i^t|_{\left(X_j^j\right)\times Y}(\sigma,\pi')=\chi_i^t|_{\left(X_j^j\right)\times Y}(\sigma',\pi')$. As firm i's beliefs about $\left(c_{j,t},s_{t-1}^j\right)_j$ under strategy profile σ with information structure π' are $\left(\left(X_{j,i}^j\right)_{j=1}^n\times Y_i^\tau\right)_{\tau=0}^{t-1}$-measurable, this implies that firm i's belief about $\left(c_{j,t},s_{t-1}^j\right)_j$ at any history $h_i^t\in H_i^t(\sigma,\pi')$ under strategy profile σ with information structure π' is a convex combination of its beliefs about $\left(c_{j,t},s_{t-1}^j\right)_j$ at various histories $h_i^t\in H_i^t(\sigma',\pi')$ under strategy profile σ' with information structure π'. Combining these observations, we see that firm i's belief about $\left(c_{j,t},s_{t-1}^j\right)_j$ at any history $h_i^t\in H_i^t(\sigma,\pi)$ under strategy profile σ with information structure π is a convex combination of its beliefs about $\left(c_{j,t-1},s_{t-1}^j\right)_j$ at various histories $h_i^t\in H_i^t(\sigma',\pi')$ under strategy profile σ' with information structure π'. Hence, $\sup_{h_i^t\in H_i^t\left(\sigma',\pi',c_{i,t}^i,s_{t-1}^i\right)}d_i^{\sigma',\pi'}(h_i^t)\geq\sup_{h_i^t\in H_i^t\left(\sigma,\pi,c_{i,t}^i,s_{t-1}^i\right)}d_i^{\sigma,\pi}(h_i^t)$ for all $\left(c_{i,t}^i,s_{t-1}^i\right)\in C_i^i\times S_i$. ■

The intuition for Proposition 1 is as follows. First, a key feature of the home market principle is that it does not require firms to have any information about their competitors' prices, costs, and sales to price optimally in their own markets (by Assumption 3). Second, such information is also not required to identify deviations, as a firm can always detect entry into its home market (by Assumption 1). On the other hand, this information is useful for predicting prices and demand in the foreign markets, which in turn gives a firm access to deviations which are better tailored to foreign market conditions—and hence are more profitable. Providing this information thus increases the discount factor required for sustaining collusion.[15]

Of course, in reality information about competitors' prices, costs, and sales can sometimes provide additional information about own-market conditions and can help identify violations of market segmentation. For example, this would occur if we dropped Assumption 3, as we do in Section 2.4. The point of Proposition 1 is simply to highlight an opposing force favoring privacy that, in the context of a particular market application, would need to be weighed

[15] The role of Assumption 2 is more subtle. Without this assumption, letting firm i enter market j with a finite but uncompetitive price p_i^j could have the advantage of obscuring information about other prices or cost or demand states. In this case, more precise monitoring of p_i^j could paradoxically help the firms by sustaining less precise monitoring of other variables.

against the well-known advantages of transparency.

We conclude this section with a comment on two of the assumptions underlying Proposition 1: the assumption that costs are observable (Assumption 4) and the assumption that firms are willing to price below cost outside their home markets to punish deviators, so long as they do not expect to make any sales at these unprofitable prices. Interestingly, Proposition 1 continues to hold if we instead make the opposite assumptions, namely that costs are completely unobservable beyond the correlation among their costs $(c_j)_j$ (and in particular cannot be inferred from information about prices or demand) and firms never post unprofitable prices.

To this end, we impose the following conditions:

Assumption 5 1. There exist functions $M^C : C \to \Delta(C)$ and $M^S : S \to \Delta(S)$ such that $M(c', s) = M^C(c') M^S(s)$ for all $(c', s) \in C \times S$.

2. $\pi^{X_j^k}\left(x_j^k | p_j^k\right) = \pi^{X_j^k}\left(x_j^k | \hat{p}_j^k\right)$ for all signals x_j^k and all prices $p_j^k, \hat{p}_j^k < \infty$.

3. $\pi^Y(y|p, c', s) = \pi^Y(y|\hat{p}, \hat{c}', s)$ for all $p, \hat{p}, c', \hat{c}'$.

Definition 1 *An equilibrium is* cautious *if it satisfies $p_i^j \geq c_{i,t}^j$ for every firm i, market j, history h_i^t (on or off path), and price p_i^j played with positive probability at history h_i^t. (Note that $c_{i,t}^j$ is known to firm i at the beginning of period t even without Assumption 4, as a firm always observes its own period $t+1$ costs at the end of period t.)*

In a cautious equilibrium, the harshest punishment for firm i involves every firm j pricing at $c_{j,t}^i$ in market i. Assumption 5 ensures that firm i does not obtain any information but $(c_{i,\tau})_{\tau=0}^t$ about how the severity of this punishment evolves over the course of the game.

Proposition 2 *Under Assumptions 1–3 and 5, for any information structure π there is a cutoff discount factor $\tilde{\delta}(\pi)$ such that the home-market principle can be sustained in a cautious equilibrium if and only if $\delta \geq \tilde{\delta}(\pi)$; furthermore, if π' is more information than π then $\tilde{\delta}(\pi') \geq \tilde{\delta}(\pi)$.*

Proof. In the definition of strategy profile σ in the proof of Proposition 1, replace the off-path threat of pricing at $c_{j,t}^j$ (in every market $j \neq i$ and period t) with the threat of pricing at $c_{i,t}^j$. Under Assumption 3, this continuation strategy holds every firm to its lowest

continuation payoff in any cautious equilibrium. Let $W_i^{\sigma,\pi}(h_i^t)$ denote firm i's continuation payoff with information structure π at history h_i^t when each of its competitors follows this strategy from period $t+1$ onward. Then, arguing as in the proof of Proposition 1, the home-market principle can be sustained in a cautious equilibrium with information structure π if and only if $\delta V_i\left(c_{i,t}^i, s_{t-1}^i\right) \geq (1-\delta)\, d_i^{\sigma,\pi}(h_i^t) + \delta W_i^{\sigma,\pi}(h_i^t)$ for all i, $c_{i,t}^i$, s_{t-1}^i, and on-path histories h_i^t consistent with $\left(c_{i,t}^i, s_{t-1}^i\right)$.

Now, under Assumption 5, $W_i^{\sigma,\pi}(h_i^t)$ depends only on $\left(\left(c_{i,\tau}\right)_{\tau=0}^t, s_{t-1}^i\right)$. To see why, recall that we assume $c_i^i \leq c_j^i$ for all $i \neq j$ and all $c \in C$, and so the harshest punishment with the causious equilibrium implies that firm j (and at least one other firm) charge the price equal to the lowest marginal cost among firms $-j$, and so firm i cannot obtain the positive profit from market $j \neq i$. Hence it suffices to infer the other firms' future costs in market i and future demand in market i. Assumption 5 ensures only $\left(c_{i,\tau}\right)_{\tau=0}^t$ is informative about the former; and Assumption 2 ensures that s_{t-1}^i is sufficient to infer the latter. In particular, $W_i^{\sigma,\pi}(h_i^t)$ does not depend on σ or π.

It therefore suffices to show that the maximum deviation gain $d_i^{\sigma,\pi}(h_i^t)$ is greater under a more informative information structure. The proof of this fact is as in Proposition 1. ■

2.3 An Example where Transparency Strictly Hinders Collusion

Proposition 1 shows that, under Assumptions 1–4, giving firms more information about their competitors' prices, costs, and sales makes sustaining first-best collusion weakly more difficult. We now present a simple example—which satisfies Assumption 1–4—where collusion is strictly more difficult to sustain when firms have perfect information than when they observe only aggregate industry demand. Note that this variation in the information structure corresponds to the cartel examples described in the introduction (and discussed further in Section 3). In particular, in this example transparency about firms' actions alone—and not necessarily also transparency about the payoff-relevant cost and demand states—hinders collusion.

The example is as follows:

- The number of firms and markets n is even and ≥ 4.
- $c_i^i = 0$ and $c_i^j = c > 0$ for all $i \neq j$.

- In each market, the period 0 demand state is s_L or s_H with equal probability, with $c < s_L < s_H$. Subsequently, the period $t+1$ state is identical to the period t state with probability ϕ and switches to the other state with probability $1-\phi$, where $1/2 < \phi < 1$. Demand is thus *positively persistent*. In addition, for k odd, the demand states in markets k and $k+1$ are *perfectly negatively correlated*, while the demand states in markets k and $k+1$ are independent of the demand states in the other markets. In particular, in every period exactly half of the markets are in each demand state.[16]

- Market demand curves are linear: $D\left(\underline{p}^j, s^j\right) = s^j - \underline{p}^j$.

- Firms observe only their own prices and sales, as well as total industry demand $\sum_j D\left(p^j, s^j\right)$.

Let $\tilde{s}_L = \phi s_L + (1-\phi)\, s_H$, $\tilde{s}_H = (1-\phi)\, s_L + \phi s_H$, and $\bar{s} = (s_H + s_L)/2$. Monopoly prices are given by $p_i^m(0, s_L) = \tilde{s}_L/2$, $p_i^m(0, s_H) = \tilde{s}_H/2$, and $p_i^m(\emptyset, \emptyset) = \bar{s}/2$. Finally, assume that $s_L > \tilde{s}_H/2$ and $\tilde{s}_L/2 > c$. Note that first-best industry profits are given by

$$n\left[(1-\delta)\frac{\bar{s}^2}{4} + \delta\left(\frac{\tilde{s}_L^2}{8} + \frac{\tilde{s}_H^2}{8}\right)\right].$$

Proposition 3 *Let δ^* (resp., δ^{**}) be the cutoff discount factors above which first-best industry profits can be sustained in equilibrium when firms observe only industry demand (resp., observe all prices and sales). Then $\delta^* < \delta^{**}$. In particular, if $\delta \in (\delta^*, \delta^{**})$ then first-best industry profits can be sustained if firms observe only industry demand, but not if firms also observe each of their competitor's prices and sales.*

The intuition is simple. Under the home-market principle, the most tempting deviation is to set the myopically optimal prices in all markets, in a period when demand in one's home market was just low. The home-market principle is therefore sustainable if and only if this deviation is deterred by the threat of reverting to zero prices in all markets. This threat is equally effective whether firms observe only industry demand or also observe the full vector of prices and quantities. However, the temptation to deviate is less when firms observe

[16] This extreme assumption on the correlation structure could easily be relaxed. As will become clear, all that is really needed is that a firm cannot determine its myopically optimal price in each market on the basis of industry demand alone.

only industry demand. In this case, a firm holds uniform beliefs over the demand state in all markets other than its home market (and the market whose demand state is perfectly negatively correlated with its home market's). In particular, the firm believes that in each of these unknown markets the home firm will price at $\tilde{s}_L/2$ or $\tilde{s}_H/2$ with equal probability. Hence, its best deviation is in every unknown market to price at either $\tilde{s}_L/2$ (winning the market for sure) or $\tilde{s}_H/2$ (winning the market with probability 1/2), for a total deviation gain of

$$(n-2)\max\left\{\frac{\tilde{s}_L\left(\tilde{s}_L-2c\right)}{4},\frac{\tilde{s}_H\left(\tilde{s}_H-2c\right)}{8}\right\}+\frac{\tilde{s}_H\left(\tilde{s}_H-2c\right)}{4},$$

where the last term is the deviation gain in the correlated market.

On the other hand, when the full vector of prices and quantities is observable, a deviator can perfectly infer the demand state in all markets, so its deviation gain becomes

$$(n-2)\left(\frac{\tilde{s}_L\left(\tilde{s}_L-2c\right)}{8}+\frac{\tilde{s}_H\left(\tilde{s}_H-2c\right)}{8}\right)+\frac{\tilde{s}_H\left(\tilde{s}_H-2c\right)}{4}.$$

Clearly, the deviation gain is strictly larger when all prices and quantities are observable. Hence, in this case firms must be strictly more patient to sustain first-best industry profits.

Two remarks on this example:

(1) In the example, there is no advantage to firms' observing industry demand rather than observing *nothing* beyond their own prices and sales. Why might it be beneficial for firms to observe industry demand? Suppose that, unlike in the example, minmaxing a deviator requires Nash reversion by *all* firms, rather than only a single firms. (For instance, this would be the case if firms have capacity constraints that lie between the monopoly and competitive quantities, so that a single firm can fulfill monopoly demand in her home market but cannot fulfill the competitive demand in any market.) Then there is a benefit to alerting all firms whenever a price cut occurs in any market. In the current setting with no aggregate demand uncertainty, a simple way of doing this is by letting the firms observe industry demand.

(2) Suppose firms can also observe each other's prices, while still observing only their own sales and industry demand. (Note that this is the version of the model with perfect monitoring of actions but no monitoring of the payoff-relevant demand state). Under the first-best action plan, a firm's price in period $t-1$ perfectly reveals her home market's

demand state in period $t-2$. Hence, a firm contemplating a deviation in period t can infer all markets' demand states in period $t-2$. If the firm prices at $\tilde{s}_L/2$ in markets with low demand at $t-2$ and prices at $\tilde{s}_H/2$ in markets with high demand at $t-2$, it receives expected payoff

$$\frac{\tilde{s}_L\left(\tilde{s}_L-2c\right)}{8}+\phi\frac{\tilde{s}_H\left(\tilde{s}_H-2c\right)}{8}.$$

A firm's maximum deviation gain in this model is therefore

$$(n-2)\max\left\{\frac{\tilde{s}_L\left(\tilde{s}_L-2c\right)}{4},\frac{\tilde{s}_H\left(\tilde{s}_H-2c\right)}{8},\frac{\tilde{s}_L\left(\tilde{s}_L-2c\right)}{8}+\phi\frac{\tilde{s}_H\left(\tilde{s}_H-2c\right)}{8}\right\}+\frac{\tilde{s}_H\left(\tilde{s}_H-2c\right)}{4}.$$

As $\phi<1$, it follows that the critical discount factor for first-best industry profits to be sustainable when only prices are observable lies in the interval $[\delta^*,\delta^{**})$. Thus, for some discount factors first-best profits are ruled out when all firms' prices are observable, even if their sales remain unknown.

2.4 An Example where the Optimal Level of Transparency Varies with the Discount Factor

In the next example, Assumption 3 is violated and the optimal level of transparency varies with the discount factor (and may be intermediate between full transparency and full privacy). Specifically, the more patient are the firms, the higher is the level of transparency that maximizes sustainable industry profits. Full transparency is thus optimal for sufficiently patient firms, but not for impatient ones.

There are two firms, two markets, and two market demand states: s_L and s_H. Period $t+1$ demand in the two markets is perfectly correlated, and is also perfectly correlated with a random variable s_t^3 that realizes in period t. (The notation indicates that this random variable can be interpreted as the period t demand state in a third, "dummy" market in which all firms have infinite production costs.) The state s_t^3 is i.i.d. across periods with $\Pr\left(s_t^3=s_L\right)=\Pr\left(s_t^3=s_H\right)=1/2$. Demand in each market $j=1,2$ is again linear: $D\left(\underline{p}^j,s^j\right)=s^j-\underline{p}^j$. Each firm can produce at zero cost in either market.[17] Assume that

[17] For this example, it would be equivalent to assume that there is only a single market and that ties are broken randomly rather than in favor of the home firm.

$\delta \geq 1/2$.

At the end of period t, the firms observe a common signal z_t of the state s_t^3, with $z_t \in \{z_L, z_H\}$ and

$$\pi\left(z_L|s_L\right) = \pi\left(z_H|s_H\right) = \phi,$$

where $\phi \geq 1/2$ now measures the informativeness of the signal. (This reuse of notation from the previous example will also let us use the same notation for $\tilde{s}_L$ and $\tilde{s}_H$.) For example, if $\phi = 1$ then at the end of period t the firms perfectly learn the period $t+1$ demand state in both markets, while if $\phi = 1/2$ they obtain no information about the period $t+1$ demand state.

This information structure violates Assumption 3, as (for instance) s_{t+1}^1 is not independent of s_t^3 conditional on s_t^1. Intuitively, making the signal z_t more informative now comes with the benefit of allowing more accurate pricing in each firm's home market, as well as the cost of allowing more accurate deviations in the foreign market. We are interested in solving for the level of informativeness ϕ that allows for the greatest industry profits and in investigating how this depends on the parameters of the model.

To do this, fix an arbitrary Markovian equilibrium and let u_L and u_H be a firm's period $t+1$ profits when z_t equals z_L and z_H, respectively.[18] As s_t—and therefore z_t—are i.i.d. across periods with equal probability on each realization, a firm's incentive compatibility constraints following signals z_L and z_H are respectively

$$\begin{aligned} \left(1-\delta\right)u_L &\leq \delta\left(\frac{u_L}{2} + \frac{u_H}{2}\right) \text{ and} \\ \left(1-\delta\right)u_H &\leq \delta\left(\frac{u_L}{2} + \frac{u_H}{2}\right). \end{aligned}$$

Consider the problem of maximizing profits $\left(u_L + u_H\right)/2$ subject to these constraints. Assuming that only the second constraint binds at the optimum for $\delta \geq \frac{1}{2}$ (as can be checked), it is optimal to set $u_L = \tilde{s}_L^2/4$, and the binding constraint becomes

$$u_H \leq \left(\frac{\delta}{2-3\delta}\right)\frac{\tilde{s}_L^2}{4}.$$

[18] A Markovian equilibrium is one in which on-path play in period $t+1$ is a function only of z_t. We restrict attention to Markovian equilibrium only in the current example.

The optimal equilibrium is therefore given by setting $u_L = \tilde{s}_L^2/4$ and

$$u_H = \min\left\{\left(\frac{\delta}{2-3\delta}\right)\frac{\tilde{s}_L^2}{4}, \frac{\tilde{s}_H^2}{4}\right\}.$$

Hence, optimal (per market) industry profits equal

$$V^* = \min\left\{\frac{1}{2}\left(1+\frac{\delta}{2-3\delta}\right)\frac{\tilde{s}_L^2}{4}, \frac{1}{2}\left(\frac{\tilde{s}_L^2}{4}+\frac{\tilde{s}_H^2}{4}\right)\right\}.$$

How does V^* vary with ϕ and δ? Let $\delta^*(\phi)$ be the value of δ that equalizes the bracketed terms. Note that $\tilde{s}_L^2$ is decreasing in ϕ while $\tilde{s}_L^2+\tilde{s}_H^2$ is increasing in ϕ (by Jensen's inequality), so $\delta^*(\phi)$ is increasing in ϕ. In addition, $\delta^*(1/2) = 1/2$, and

$$\delta^*(1) = \frac{2}{3+\left(\frac{s_L}{s_H}\right)^2}.$$

There are two cases:

1. If $\delta < \delta^*(\phi)$, then $V^* = \frac{1}{2}\left(1+\frac{\delta}{2-3\delta}\right)\frac{\tilde{s}_L^2}{4}$, and V^* is increasing in δ and decreasing in ϕ.
2. If $\delta > \delta^*(\phi)$, then $V^* = \frac{1}{2}\left(\frac{\tilde{s}_L^2}{4}+\frac{\tilde{s}_H^2}{4}\right)$, and V^* is constant in δ and increasing in ϕ.

We can now read off the optimal value of ϕ. Let $\phi^*(\delta) = \{\phi : \delta^*(\phi) = \delta\}$, and note that $\phi^*(\delta)$ is an increasing function on $\delta \in [1/2, \delta^*(1)]$.

Proposition 4 *The level of informativeness ϕ that maximizes industry profits V^* is given by $\phi = \min\{\phi^*(\delta), 1\}$.*

Proof. Follows because V^* is increasing in ϕ if $\phi < \phi^*(\delta)$ and decreasing in ϕ if $\phi > \phi^*(\delta)$. ■

Thus, the optimal level of informativeness is increasing in δ, and it lies strictly between 0 and 1 when $\delta \in (1/2, \delta^*(1))$.[19]

To understand this result, note that giving the firms more information about the state always increases unconditional expected first-best profits, $(\tilde{s}_L^2 + \tilde{s}_H^2)/8$, but decreases expected

[19] The analysis here would also be exactly the same if demand in the two markets were perfectly negatively correlated, rather than positively correlated. The only difference is that incentive compatibility would bind for one firm after each signal, rather than binding for both firms after the high signal.

first-best profits after the bad signal, $\tilde{s}_L^2/4$. If δ is close to 1/2, then incentive compatibility implies that profits in the low and high states cannot be too different, which implies that profits in both states must be close to $\tilde{s}_L^2/4$.[20] Therefore, if δ is close to 1/2 then the optimal information structure is less informative (to maximize $\tilde{s}_L^2/4$), while if δ is close to 1 then the optimal information structure is more informative (to maximize $\left(\tilde{s}_L^2 + \tilde{s}_H^2\right)/8$).[21]

3 Coarse Information and the Home-Market Principle in European Industrial Cartels

The inner workings of real-world cartels are inevitably far more complicated and nuanced than those of any theoretical model. Nonetheless, we believe that some of the key mechanisms underlying several recent major industrial cartels are quite consistent with the results presented thus far. In this sense, our results may be viewed as one possible explanation for some aspects of the behavior of these cartels, especially the puzzling efforts on the part of some cartels to maintain the privacy of their participants' firm-level data in support of the home-market principle. Of course, for all of the cartels discussed in this section, our model does not provide the only possible explanation for the documented behavior, and our aim is not to rule out other explanations. The goal is only to connect our theory with some observed cartels and to highlight the theory's plausibility.

The discussion of the following cases is based on antitrust decisions of the European Commission (EC), as well as on summaries and analysis of these decisions by Harrington (2006) and Marshall and Marx (2012).

Copper Plumbing Tubes: "Copper plumbing tubes are used for water, oil, gas and heating installations in the construction industry. The main customers are distributors, wholesalers, and retailers that sell the plumbing tubes to installers and other end consumers," (Harrington, p. 85; EC—Copper Plumbing Tubes, pp. 10–11). From 1988 to 2001, the European copper plumbing tubes industry—a roughly 1 billion Euro industry—was cartelized

[20] This is also an implication of Proposition 4 of Kandori (1991b).

[21] The same argument implies that, in standard price competition models with time-varying demand—such as Rotemberg and Saloner (1986), Haltiwanger and Harrington (1991), Kandori (1991b), and Bagwell and Staiger (1997)—collusion can be easier to sustain if the firms do not observe the current demand state. Hochman and Segev (2010) derive a similar result in a model of repeated international trade policy.

by a group of five to nine firms that jointly accounted for approximately 65% to 80% of the market (EC, pp. 15–16). The operation of the European copper plumbing tubes cartel reflects several key features of our model.

First, as the cartel grew and intensified, it developed an increasingly sophisticated and formalized approach to information-sharing among its members. Over time, the cartel shifted from informally self-reporting prices and sales, to reporting to a trade association—the International Wrought Copper Council (IWCC)—and finally to reporting to the World Bureau of Metal Statistics (WBMS), a statistical bureau. The WBMS was eventually so linked to the cartel that the EC viewed providing information to the WBMS as prima facie evidence of participation in the cartel, ruling that in the case of one firm ("Halcor") "Halcor's continued supply of sales volumes to the WBMS can only be understood as meaning that Halcor had not taken a final decision to completely withdraw entirely from the illegal arrangements," (p. 129).

What is most interesting from the perspective of our model is that—despite being engaged in clearly illegal activities—the cartel participants seemed to exchange *less* detailed information as their means of exchanging information improved. In the early period in which the cartel relied on informal self-reporting, "Each producer provided Mr. [...] with its volume figures of deliveries on a country-by-country basis on a monthly or quarterly basis. With these figures, Mr. [...] prepared a "spreadsheet" that contained the collected data," (p. 57). But, later on, "As of 1 January 1998, a data exchange took place initially on a monthly, later on a quarterly basis through the [WBMS]. WBMS statistics only contained aggregated figures and no company specific information," (p. 52) with the aim of "enabling each individual participant to calculate his share of the business as a percentage of the total business of the participants," (p. 75). It thus appears that the cartel shifted from sharing firm-level data to aggregate data as cooperation within the cartel intensified.

Consistent with its reliance on coarse information, the copper plumbing tubes cartel operated on the home-market principle. The EC ruled that "...the basic goal of the [cartel] meetings was to protect the main producers' home markets and to freeze the market shares..." (p. 57). In addition, as in our model, designated "market leaders" were responsible for setting prices and monitoring adherence to the collusive agreement within their home markets: "Indeed, part of the arrangements concerned the organisation of a mechanism of

market segregation: national markets were given a market leader who would decide the price variations," (p. 169). In summarizing its decision, the EC wrote that the cartel "ensured implementation of the market allocation and price agreements/coordination by a monitoring system consisting of a market leader arrangement for various European territories," (p. 115).

The European copper plumbing tubes cartel was thus based on the home-market principle and sustained collusion by exchanging only aggregate data, despite its apparent ability to exchange more detailed information. This combination of features is consistent with the predictions of our model.

Isostatic Graphite: Isostatic graphite is a graphite product used in industrial applications such as the production of certain types of electrodes and semiconductors (EC—Specialty Graphite, p. 6). The EC prosecuted eight firms for cartelizing the European isostatic graphite industry (roughly a 500 million Euro industry) in the mid-1990's. The cartel operated through meetings at both the European and country levels. The striking example of using a calculator to keep firm-level sales secret comes from the Italian country-level meetings. According to the EC, "A common practice in the meeting... consisted in trying to determine the size of the market by passing around a calculator where each participant entered its company's sales volumes of isostatic products. This ensured that no one saw the individual companies' volumes, but only aggregate sales to the Italian market," (p. 61). The isostatic graphite cartel also relied on the home-market principle, fixing national market shares in the European meetings and dividing up large customers in the country-level meetings: "in particular at local level, the exchanges of information concerned the repartition of major customers," (p. 25). In the Italian market, "a list of sixteen major customers was prepared and it was agreed to freeze the respective sales shares for them," (p. 63). It thus appears that—at least in the Italian market—the isostatic graphite cartel also relied on a combination of the home-market principle and the deliberate coarsening of exchanged information.

Other cartels: While the wealth of institutional detail surrounding major cartels can make it hard to pinpoint the exact mechanisms used to support collusion, references to information coarsening and (especially) the home-market principle are common in the EC decisions. The home-market principle (implemented through either exclusive territories or the allocation of individual large customers) was the basis of the cartels in choline chlo-

ride, district heating pipes, electrical and mechanical carbon graphite, lysine, methionine, nucleotides, seamless steel tubes, soda ash, vitamins, and zinc phosphate (Harrington, pp. 34–40), in addition to the copper plumbing tubes and isostatic graphite cartels already discussed. Information coarsening—in particular, the practice of firms' reporting detailed individual-level data to an intermediary, which then returned only aggregate data to the firms—also seems have played an important role in the cartels in plasterboard (Harrington, p. 54) and low density polyethylene (Marshall and Marx, p. 132). For example, in the plasterboard cartel, "Four firms set up a system for exchanging information through an independent expert, Mr. [U, independent consultant]. The operation was placed under the aegis of the Plasterboard Industry Group. Each producer gave its figure to Mr. [U] on a confidential basis and the results were compiled in the latter's office, giving an aggregate figure, which was then sent to the participants. This figure enabled each producer to calculate its own market share, but not that of the others," (EU—Plasterboard, p. 54). The plasterboard cartel also seemed to rely to some extent on the home-market principle: in the EC's summary of the cartel's infringement, it found that the participants had "a view to sharing out or at least stabilising the German market," (p. 6).

In sum, both information coarsening and the home-market principle appear to have been important features of several major European industrial cartels.

4 Transparency can Hinder Collusion in Repeated Games: Optimally Rocky Relationships

In the model considered so far, optimal collusive equilibria were "rocky"—that is, stochastic—for the obvious reason that the underlying firm costs and market demand states were assumed to be stochastic. As we have seen, this stochasticity opened up the possibility that a transparent information structure could give too much information to potential deviators and thereby hinder collusion. In reality, there could be a variety of reasons why a cartel would like to rely on a stochastic collusive agreement: for example, such an agreement could be harder for antitrust authorities to detect or could be less susceptible to attack by outside entrants. In any such situation, the mechanism identified in Proposition 1 will encourage

the cartel to adopt a non-transparent information structure.

In this section, we probe the limits of this theoretical mechanism by asking whether it can also be optimal for a cartel to follow a stochastic equilibrium even if the physical environment is completely stationary and free of exogenous factors favoring stochasticity. Perhaps surprisingly, we show that the answer is yes. The key idea is that—even in a stationary physical environment—the need to provide intertemporal incentives alone can lead the cartel to follow a stochastic equilibrium. This again opens up the possibility that transparency can hinder collusion. Specifically, we exhibit an example of a repeated game in which higher joint profits can be sustained when firms observe only their own actions and profits than when they also observe their competitors' actions.

Our example is the following canonical normal-form entry deterrence game:

$$\begin{array}{rcc} & Out & In \\ Fight & 3,0 & -3,-3 \\ Accommodate & 4,0 & 0,3 \end{array} \tag{1}$$

While the entry deterrence game is not an obvious model of collusion, this game can easily arise within a cartel that is attempting to implement the home-market principle: Suppose there are two firms and two markets, and that firm 1 is unable to enter firm 2's market while firm 2 is able to enter firm 1's market but has a cost disadvantage in this market. (For example, a tariff or technological barrier may make it unprofitable for firm 1 to enter firm 2's market, while no such exogenous protection applies to firm 1's market.) Suppose further that firm 1 can either price low to deter entry by firm 2 ($Fight$) or price high to extract surplus from the consumers in her market ($Accommodate$). In this setting, firm 2 will always obtain monopoly profits in its home market (which can thus be ignored in the analysis), and the remaining game in firm 1's market corresponds to the entry deterrence game. With this interpretation, the problem of maximizing cartel profits across the two markets reduces to maximizing the sum of payoffs in the entry deterrence game.

We will show that, when the firms' actions are perfectly observable and the discount factor equals 0.21, the only repeated game equilibrium is the infinite repetition of the static Nash equilibrium $(Accommodate, In)$, yielding a joint profit of 3. This remains true when

the firms can use a mediator to help them correlate their actions. On the other hand, when each firm observes only its own action and profit, we will show that with the assistance of a mediator the firms can attain a joint profit of at least 3.5.

Before stating the result formally, let us comment on the assumption that the firms have access to a mediator. This assumption seems quite realistic in light of the examples discussed above: many real-world cartels do rely on intermediaries to help them collude, and one of the intermediary's roles is often summarizing (and thus coarsening) more detailed information about the cartel participants' behavior.[22] Nonetheless, an alternative interpretation is that the mediator is simply a stand-in for the various imperfect private monitoring structures under which the firms could interact: with this interpretation, a "message" of the mediator's would instead be interpreted as a firm's signal of its competitors' behavior. In Sugaya and Wolitzky (2016), we discuss what properties such a private monitoring structure would need to have to justify this interpretation.[23]

4.1 Impossibility of Entry Deterrence with Mediated Perfect Monitoring

We first show that non-Nash outcomes of the entry deterrence game cannot be sustained under perfect monitoring, even if the firms have access to a mediator who also perfectly observes their actions.[24] A standard application of the revelation principle (Forges, 1986; Myerson, 1986) implies that in such games there is no loss of generality in restricting attention to so-called obedient equilibria, where in each period the game proceeds as follows:[25]

1. The mediator makes a private action recommendation $r_{i,t} \in A_i$ to each firm $i = 1, 2$. This recommendation may be conditioned on the mediator's private history, $h^t =$

[22] Levenstein and Suslow (2006, p. 69) report that "Industry associations often engage in the collection and dissemination of information, which may facilitate collusion. Between a quarter and a half of the cartels in U.S. cross-section studies report the involvement of trade associations."

[23] There is also a methodological reason for allowing a mediator: without a mediator, firms could benefit from observing noisier signals simply because this gives them new ways to correlate their actions, rather than because they benefit from the lack of transparency per se.

[24] When we consider imperfect monitoring in Section 4.2, we will ultimately require the mediator to rely on the firms' reports of their own actions. If anything, this difference in the mediator's information structure stacks the deck in favor of perfect monitoring.

[25] For a more detailed exposition of repeated games with mediated perfect monitoring, see Sugaya and Wolitzky (2016).

$(r_\tau, a_\tau)_{\tau=0}^{t-1}$, which consists of the entire history of both private recommendation profiles and public action profiles.

2. Each firm i takes an action $a_{i,t} \in A_i$. This action may be conditioned on firm i's private history, $h_i^t = (r_{i,\tau}, a_\tau)_{\tau=0}^{t-1}$, which consists of the history of private recommendations to firm i and public action profiles. By definition of an obedient equilibrium, firm i obeys the mediator's recommendation (i.e., plays $a_{i,t} = r_{i,t}$) on the equilibrium path.

3. Both the mediator and the firms observe the realized action profile a_t.

Proposition 5 *In the entry deterrence game with $\delta = 0.21$, the only sequential equilibrium with mediated perfect monitoring is the infinite repetition of $(Accommodate, In)$, which yields a joint profit of 3.*

Proof. We show that, in every sequential equilibrium, firm 1 plays *Accommodate* at every history. Given this, the only equilibrium is the infinite repetition of $(Accommodate, In)$.

Suppose firm 1 plays $Fight$ with positive probability at some period t history. When recommended $Fight$, firm 1 can gain instantaneous utility 1 by deviating. Hence, for $Fight$ to be incentive compatible, $(Accommodate, Out)$ must then be played in period $t+1$ with probability at least x such that

$$(1-\delta) \underbrace{(1)}_{\text{deviation gain}} + \delta \underbrace{(0)}_{\text{minmax payoff}} = \delta(1-\delta) \underbrace{[x(4) + (1-x)(3)]}_{\text{maximum period } t+1 \text{ payoff}} + \delta^2 \underbrace{(4)}_{\text{maximum payoff}},$$

or

$$x = \frac{1 - 4\delta - \delta^2}{\delta(1-\delta)} \approx 0.699 > 0.65.$$

Therefore, when $Fight$ is played in period t and firm 2 is then recommended Out in period $t+1$, firm 2 believes that firm 1 will play *Accommodate* with probability at least 0.65. However, for Out to be incentive compatible for firm 2, this probability cannot exceed y such that

$$(1-\delta) \underbrace{[y(3) + (1-y)(-3)]}_{\text{deviation gain}} + \delta \underbrace{(0)}_{\text{minmax payoff}} = \delta \underbrace{(3)}_{\text{maximum payoff}},$$

or

$$y = \frac{1}{2(1-\delta)} \approx 0.633.$$

This is a contradiction. ■

4.2 Possibility of Entry Deterrence with Mediated Private Monitoring

We now construct a sequential equilibrium with mediated private monitoring with joint profit 3.5. Given Proposition 5, this shows that transparency (i.e., perfect monitoring of actions) can hinder entry deterrence, and thus collusion.

In this section, we assume that in every period each firm i observe its own action a_i and profit u_i, but not its competitor's action a_{-i}. Meanwhile, the mediator observes the entire realized action profile $a = (a_1, a_2)$. We return to the assumption that the mediator is in this sense "omniscient" at the end of the section.

Formally, in each period the game is as follows:

1. The mediator makes a private action recommendation $r_{i,t} \in A_i$ to each firm i. This recommendation may be conditioned on the mediator's private history $h^t = (r_\tau, a_\tau)_{\tau=0}^{t-1}$.

2. Each firm i takes an action $a_{i,t} \in A_i$. This action may be conditioned on firm i's private history, $h_i^t = (r_{i,\tau}, a_{i,\tau}, u_{i,\tau})_{\tau=0}^{t-1}$. (Note the distinction with the perfect monitoring case, where firm i can also condition its action on $(a_{-i,\tau})_{\tau=0}^{t-1}$.)

3. Firm i observes its payoff $u_{i,t} = u_i(a_t)$. The mediator observes the realized action profile a_t.

Compared to the mediated perfect monitoring case, mediated private monitoring offers *less feedback* to each firm. Imagine that the mediator is a dynamic mechanism designer who wants to implement an efficient action path while facing incentive compatibility constraints for each firm. Giving a firm more information splits its information sets and therefore tightens the incentive compatibility constraints. An omniscient mediator has the same information as the mediator with perfect monitoring, while the firms have less information under mediated private monitoring than under mediated perfect monitoring. This implies that the equilibrium payoff set with mediated private monitoring (with an omniscient mediator) is

always at least weakly larger than that with mediated perfect monitoring. Our contribution here is to show that it is strictly larger in the entry deterrence game.[26]

Proposition 6 *In the entry deterrence game with $\delta = 0.21$, there exists a sequential equilibrium with mediated private monitoring which yields a joint profit of* 3.5.

Proof. Consider the following strategy for the mediator, which yields joint profit 3.5 when the firms follow their recommendations:

- To determine firm 1's recommendation, flip a fair coin at the beginning of the game. If it comes up heads, recommend $Fight$ in odd periods and $Accommodate$ in even periods. If it comes up tails, recommend $Fight$ in even periods and $Accommodate$ in odd periods. Crucially, do not disclose the result of the coin toss to firm 2.
- Always recommend Out to firm 2.
- If firm 1 ever unilaterally disobeys its recommendation, recommend $(Accommodate, In)$ forever.
- If firm 2 ever unilaterally disobeys its recommendation, restart the strategy profile, flipping a fresh coin to determine firm 1's recommendations.

To see that this is an equilibrium, first note that playing Out is incentive compatible for firm 2 because firm 2 always assigns equal probability to $Fight$ and $Accommodate$ (so both firm 2's instantaneous payoff and continuation payoff from both Out and In are always equal to zero).[27] Finally, the non-trivial incentive constraint for firm 1 comes when she is recommended $Fight$, and is given by

$$(1-\delta)(1)+\delta(0) \leq \frac{\delta}{1+\delta}(4)+\frac{\delta^2}{1+\delta}(3),$$

[26] This discussion also implies the following *characterization* of when transparency hinders collusion: Fix an equilibrium with uninformed firms and an omniscient mediator that maximizes industry profits. If, for every such equilibrium, revealing all firms' past actions leads to the violation of an incentive constraint, then maximimum industry profits are strictly higher with mediated private monitoring than with mediated perfect monitoring. Of course, the problem with this characterization is that determining the optimal equilibrium with uninformed firms and mediated private monitoring is typically intractable. This is why we content ourselves with showing that transparency can hinder collusion by means of an example.

[27] Unless the mediator has switched to recommending $(Accommodate, In)$, in which case In is clearly incentive compatible.

or

$$\delta > \frac{1}{2}\left(\sqrt{2}-1\right) \approx 0.207.$$

As $\delta = 0.21$, we have an equilibrium. ■

What is the intuition for why transparency prevents collusion in this example? The key observation is that firm 1 is so impatient that she is willing to play $Fight$ only if she is rewarded with the profile $(Accommodate, Out)$ with high probability in the *very next period*. With perfect monitoring, this means that, when firm 1 plays $Fight$ in period t, firm 2 observes this action and then knows to expect $Accommodate$ with high probability in period $t+1$. But firm 2 is too impatient to play Out when firm 1 is playing $Accommodate$ with high probability, so he will deviate to In. This implies that firm 1 will never play $Fight$. Hence, inefficient entry cannot be deterred under perfect monitoring.

With private monitoring, however, the mediator can recommend that firm 1 mixes between $Fight$ and $Accommodate$ without informing firm 2 of the outcome of this randomization. Firm 2 therefore never expects firm 1 to play $Accommodate$ with probability greater than 50% and is thus willing to play Out. (And firm 2's payoff from Out is zero regardless of firm 1's action, so firm 2 does not learn anything from observing his own payoffs.) Meanwhile, as firm 1 always plays $Accommodate$ in period $t+1$ after playing $Fight$ in period t, she receives the reward of $(Accommodate, Out)$ required to make $Fight$ incentive compatible. This arrangement therefore succeeds in deterring entry.

Concisely put, in this example collusion requires a stochastic intertemporal dependence between plays of $Fight$ and $Accommodate$, and along such a rocky path of play revealing firm 1's past actions will prompt a deviation by firm 2.

Finally, thus far we have simplified the exposition by assuming that the mediator directly observes the firms' actions and in particular does not need to rely on self-reports by the firms. This assumption may be realistic: according to Harrington (2006), the industry groups and accounting firms supporting the vitamins, plasterboard, and citric acid cartels directly audited cartel participants to make sure they were reporting their sales truthfully. Nonetheless, we show in Appendix B that the assumption that the mediator is omniscient can be completely dispensed with if firm profits are noisy and the game is augmented with a mutual minmax Nash equilibrium action profile. This augmentation of the game—corresponding

for example to the possibility of a mutually destructive price war in firm 1's market—is necessary for the construction because, when the mediator must rely on self-reports by the firms to detect a deviation, it will not be able to distinguish between misreports by firm 1 and misreports by firm 2, so there must be a single action profile that can be used to punish both firms simultaneously.

5 Transparency Helps Collusion in Concave Repeated Games: Optimally Smooth Relationships

Our analysis so far has been concerned with deriving conditions under which firms can benefit from maintaining privacy within a cartel. To better understand the boundaries of these results, we now turn to the converse question of whether there exist natural classes of games in which the conventional wisdom that transparency can only help support collusion is borne out. In contrast to our results so far, we show that the answer is yes. In particular, we show that transparency always favors collusion in repeated (non-stochastic) games when the stage game is in the following class.

Definition 2 *A stage game $G = (N, A, u)$ is* concave *if the following conditions hold for all $i \in N$:*

- *A_i is a compact and convex subset of $\mathbb{R}$.*
- *$u_i(a)$ is continuous and jointly concave in the action profile a.*
- *Defining firm i's* deviation gain *from a recommended action $a_i \in A_i$ against a mixed action profile of her opponents' $\alpha_{-i} \in \Delta(A_{-i})$ by*

$$d_i(a_i, \alpha_{-i}) = \max_{\tilde{a}_i \in A_i} \mathbb{E}^{\alpha_{-i}} \left[u_i(\tilde{a}_i, a_{-i}) - u_i(a) \right],$$

there exist functions $f : \mathbb{R}^2 \to \mathbb{R}$ and $g : \mathbb{R}^{n-1} \to \mathbb{R}$ such that

1. *$d_i(a_i, \alpha_{-i}) = f\left(a_i, \mathbb{E}^{\alpha_{-i}}\left[g(a_{-i})\right]\right)$,*
2. *f is convex, and*

3. either

(a) $g(a_{-i}) = \sum_{j \neq i} w_j a_j$ *for weights* $w_j \in \mathbb{R}$*, or*

(b) f *is increasing in its second argument and* g *is convex.*

We show below that this class of games includes n-player linear Cournot and differentiated-product Bertrand competition.[28]

The economic content of the definition is that, in a concave game, replacing a distribution of actions with its expectation increases players' payoffs (part 2 of the definition) while reducing their incentives to deviate (part 3). This implies that optimal equilibria in concave games have deterministic equilibrium paths of play. In addition, if an equilibrium has a deterministic path of play, then transparency gives deviators no actionable information. Hence, transparency is the optimal information structure in a concave game.

Denote firm i's correlated minmax payoff by

$$\underline{u}_i = \min_{\alpha_{-i} \in \Delta(A_{-i})} \max_{a_i \in A_i} u_i(a_i, \alpha_{-i}).$$

We show that, if it is possible to hold a deviant firm to its minmax payoff, full transparency always favors collusion in the sense of letting the firms sustain higher profits.

Proposition 7 *Suppose the stage game is concave, and suppose that, with perfect monitoring, for each firm* i *there exists a subgame perfect equilibrium that gives firm* i *payoff* $\underline{u}_i$*. Then, for any monitoring structure (perfect or imperfect, with or without a mediator) and any payoff vector* $v \in \mathbb{R}^n$ *sustainable in Nash equilibrium under that monitoring structure, there exists a subgame perfect equilibrium under perfect monitoring that Pareto dominates* v*.*

Proof. Fix a monitoring structure and a Nash equilibrium. Letting $\mathbb{E}[\cdot]$ denote ex ante expectation as taken at the beginning of the game under this equilibrium, consider the following strategy profile under perfect monitoring:

- On path, player i plays $\mathbb{E}[a_{i,t}]$ in period t.[29]

[28] Under linear differentiated-product Bertrand competition, $u_i(a)$ is not concave, but $\sum_i u_i(a)$ is. Proposition 8 below shows that this is sufficient for full transparency to maximize industry profits.

[29] Note that $\mathbb{E}[a_{i,t}] \in A_i$ because A_i is convex.

- If player i unilaterally deviates (on or off path), switch to a subgame perfect equilibrium that gives her payoff $\underline{u}_i$. Ignore simultaneous deviations.

Since u_i is concave for all i, this strategy profile Pareto dominates the original equilibrium. We show that it is a subgame perfect equilibrium.

Under the original equilibrium, for every player i and every on-path period t private history h_i^t, we have

$$d_i\left(a_{i,t}, \alpha_{-i,t}|_{h_i^t}\right) \leq \sum_{s=1}^{\infty} \delta^s \mathbb{E}\left[u_i\left(a_{s+t}\right) - \underline{u} | h_i^t\right], \tag{2}$$

where $\alpha_{-i}|_{h_i^t}$ denotes the distribution over player $-i$'s actions conditional on reaching history h_i^t. Note that

$$\mathbb{E}\left[\mathbb{E}\left[u_i\left(a_{s+t}\right)|h_i^t\right]\right] \leq u_i\left(\mathbb{E}\left[a_{s+t}\right]\right),$$

by concavity of u_i and the law of iterated expectation, and

$$\begin{aligned}
\mathbb{E}\left[d_i\left(a_{i,t}, \alpha_{-i,t}|_{h_i^t}\right)\right] &= \mathbb{E}\left[f\left(a_{i,t}, \mathbb{E}^{\alpha_{-i,t}|_{h_i^t}}\left[g\left(a_{-i,t}\right)\right]\right)\right] \\
&\geq f\left(\mathbb{E}\left[a_{i,t}\right], \mathbb{E}\left[g\left(a_{-i,t}\right)\right]\right) \\
&\quad \text{(by convexity of } f \text{ and the law of iterated expectation)} \\
&\geq f\left(\mathbb{E}\left[a_{i,t}\right], g\left(\mathbb{E}\left[a_{-i,t}\right]\right)\right) \\
&\quad \text{(by condition a. or b. in Definition 2).}
\end{aligned}$$

Hence, taking expectations of both sides of (2), we have

$$f\left(\mathbb{E}\left[a_{i,t}\right], g\left(\mathbb{E}\left[a_{-i,t}\right]\right)\right) \leq \sum_{s=1}^{\infty} \delta^s \left(u_i\left(\mathbb{E}\left[a_{s+t}\right]\right) - \underline{u}\right) \text{ for all } i \text{ and } t.$$

But this is precisely the condition for the constructed strategy profile to be a subgame perfect equilibrium under perfect monitoring. ■

If in addition the stage game is symmetric and our concern is maximizing total industry profits, then we can establish the stronger result that a *stationary* equilibrium is optimal. It also suffices for this result to assume that only total industry profits $V(a) = \sum_i u_i(a)$ are concave in actions, rather than assuming that $u_i(a)$ is concave for all i; this modification

allows the current result to cover linear differentiated-product Bertrand competition. Let

$$A^* = \left\{ a \in A : \sum_i d_i(a) \leq \frac{\delta}{1-\delta} \left[V(a) - n\underline{u} \right] \right\}.$$

Thus, A^* is the set of pure action profiles a such that, when a is to be played forever in equilibrium, the sum of the firms' deviation gains at a does not exceed the present value of the continuation value that would be lost from permanently switching to a mutual minmax profile. Note that A^* is a convex, closed, and symmetric set, as A is a convex, closed, and symmetric set, V is concave, and f is convex and satisfies condition a. or b. in Definition 2. Let $V^* = \max_{a \in A^*} V(a)$, and let $a^* \in A^*$ be a symmetric action profile with $V(a^*) = V^*$.

Proposition 8 *Suppose the stage game is symmetric and concave, with the requirement that $u_i(a)$ is concave for all i replaced by the weaker requirement that $V(a)$ is concave.*

1. *For any monitoring structure, expected per-period industry profits do not exceed V^* in any Nash equilibrium.*

2. *With perfect monitoring, if there exists a subgame perfect equilibrium that gives a firm payoff $\underline{u}$, then the infinite repetition of a^* is a subgame perfect equilibrium outcome yielding per-period industry profits V^*.*

Proof. Consider the problem of maximizing expected per-period industry profits $(1-\delta)\sum_{t=0}^{\infty} \delta^t \mathbb{E}\left[V(a_t)\right]$ over all monitoring structures and all Nash equilibria. Summing (2) over players and taking expectations, a necessary condition for a strategy profile to be a Nash equilibrium is

$$\sum_i f\left(\mathbb{E}\left[a_{i,t}\right], g\left(\mathbb{E}\left[a_{-i,t}\right]\right)\right) \leq \sum_{s=1}^{\infty} \delta^s \mathbb{E}\left[V(a_{s+t}) - n\underline{u}\right] \text{ for all } t. \tag{3}$$

Therefore, a relaxed version of this problem is

$$\max (1-\delta) \sum_{t=0}^{\infty} \delta^t \mathbb{E}\left[V(a_t)\right]$$

subject to (3). Call this relaxed problem Problem 1. Note that Problem 1 has a solution by continuity and compactness.

We claim that for any (possibly non-concave) stage game G there exists a stationary solution to Problem 1. To see this, let $V = \sup_t \mathbb{E}[V(a_t)]$. Then any solution to Problem 1 must satisfy

$$\sum_i f\left(\mathbb{E}[a_{i,t}], g\left(\mathbb{E}[a_{-i,t}]\right)\right) \leq \frac{\delta}{1-\delta}\left[V - n\underline{u}\right] \text{ for all } t.$$

Hence, if we take any solution to Problem 1 and, for every t, replace α_t with $\alpha_{\tilde{t}}$ for some $\tilde{t} \in \text{argmax}_{\tilde{t}} \mathbb{E}[V(a_{\tilde{t}})]$, the resulting stationary mixed action path is another solution to Problem 1.

Hence, the value of Problem 1 is the same as the value of the problem

$$\max_{\alpha \in \Delta(A)} \mathbb{E}^{\alpha}[V(a)]$$

subject to

$$\sum_i f\left(\mathbb{E}^{\alpha}[a_i], g\left(\mathbb{E}^{\alpha}[a_{-i}]\right)\right) \leq \frac{\delta}{1-\delta}\left[\mathbb{E}^{\alpha}[V(a)] - n\underline{u}\right].$$

Call this Problem 2.

Next, since A is convex and V is concave, there exists a deterministic solution to Problem 2. This follows because, under these assumptions, replacing α with its expectation increases the objective and relaxes the constraint.

Finally, when one restricts attention to deterministic profiles α, Problem 2 becomes exactly the program defining V^*. This shows that equilibrium industry profits cannot exceed V^* for any monitoring structure.

To complete the proof, we claim that infinite repetition of a^*, supported by the threat of transitioning to a subgame perfect equilibrium that minmaxes the deviator, is a subgame perfect equilibrium under perfect monitoring. As a^* is a symmetric action profile, in this strategy profile each firm's incentive compatibility constraint is

$$d_i(a^*) \leq \frac{\delta}{1-\delta}\left[\frac{1}{n}V(a^*) - \underline{u}\right].$$

Multiplying both sides by n, this is equivalent to the condition that $a^* \in A^*$. ■

An example of a game satisfying the conditions of Proposition 8 is n-player linear Cournot oligopoly, where the market price at quantity vector q is given by

$$\max\left\{a-\sum_j q_j, 0\right\}$$

for a constant $a>0$, so profits are given by

$$u_i(q)=\max\left\{\left(a-\sum_j q_j\right)q_i, 0\right\}.$$

Restricting attention to equilibria with positive prices along the equilibrium path, $u_i(q)$ is a concave function of q. Next, letting $\tilde{q}_{-i}$ denote a distribution over q_{-i}, we have

$$d(q_i,\tilde{q}_{-i})=\frac{1}{4}\left(a-\mathbb{E}^{\tilde{q}_{-i}}\left[\sum_{j\neq i}q_j\right]\right)^2-\left(a-\mathbb{E}^{\tilde{q}_{-i}}\left[\sum_{j\neq i}q_j\right]-q_i\right)q_i.$$

Thus, letting $g(q_{-i})=\sum_{j\neq i}q_j$ and letting

$$f(q_i,x)=\frac{1}{4}(a-x)^2-(a-x-q_i)q_i,$$

we have $d_i(q_i,\tilde{q}_{-i})=f\left(q_i,\mathbb{E}^{\tilde{q}_{-i}}[g(q_{-i})]\right)$. The eigenvalues of the Hessian matrix of $f(q_i,x)$ are 0 and 5/2, so $f(q_i,x)$ is convex. The game is therefore concave. Finally, in the current formulation without production costs, the game also admits a mutual minmax Nash equilibrium, so there is always a subgame perfect equilibrium that minmaxes a deviator.[30]

A second example is given by n-player linear differentiated-product Bertrand oligopoly. We verify that this game satisfies the conditions of Proposition 8 in Appendix C.

As a final example, we describe why the entry deterrence game in Section 4 is not concave. The simplest condition of Definition 2 it violates is that action sets are not convex subsets of $\mathbb{R}$. In particular, suppose we add a third action, $.5Fight+.5Accommodate$, for firm 1,

[30] In the more realistic version with production costs, the game does not admit a mutual minmax static Nash equilibrium, but it nonetheless admits a subgame perfect equilibrium yielding payoff $\underline{u}$ if δ is sufficiently high. The argument is as in the Bertrand case (or see Abreu, 1986).

with payoffs given by interpolating between $Fight$ and $Accommodate$:

	Out	In
$Fight$	$3, 0$	$-3, -3$
$.5Fight + .5Accommodate$	$3.5, 0$	$-1.5, 0$
$Accommodate$	$4, 0$	$0, 3$

The infinite repetition of $(.5Fight + .5Accommodate, Out)$ could then be sustained in equilibrium with perfect monitoring, yielding a joint profit of 3.5.

The argument behind Proposition 8 can also be used to the show that, in symmetric, concave games where the worst subgame perfect equilibrium payoff is something other than $\underline{u}$, the equilibrium that maximizes industry profits with perfect monitoring is stationary. This resolves an issue in the theoretical literature on collusion with perfect monitoring following Abreu (1986) (e.g., Lambson, 1987; Wernerfelt, 1989; Chang, 1991; Ross, 1992; Häckner, 1996; Compte, Jeny, and Rey, 2002), which, while focusing on computing optimal punishment paths under various assumptions, typically assumes without formal justification that the cartel's goal is sustaining the best possible *stationary* path of play. At least for the linear demand curves typically considered in the literature, Proposition 8 provides a justification for this approach.

Finally, we briefly comment on the relationship between Proposition 7 and Theorem 1 of our companion paper, Sugaya and Wolitzky (2016). Theorem 1 of that paper shows that, in two-player games, there exists a critical discount factor δ^* such that if $\delta > \delta^*$ then the sequential equilibrium payoff set with mediated perfect monitoring dominates the sequential equilibrium payoff set with any monitoring structure for all non-negative Pareto weights. Thus, that result applies to *all*, *two-player* games, *with* a mediator, and with $\delta > \delta^*$; while Proposition 7 applies to *concave*, *n-player* games, *with and without* a mediator, for *all* discount factors. The contents of the two results are therefore quite different, and the intuition and proof techniques behind them are even more dissimilar.

6 Conclusion

The goal of this paper has been a reassessment of Stigler's path-breaking idea that transparency within a cartel facilitates collusion. In contrast to this idea, we find that—under fairly general conditions—transparency hinders collusion when the cartel's objective is to segment the market according to the home-market principle. Consistent with our model, several recent European industrial cartels that operated under the home-market principle appear to have gone out of their way to preserve the privacy of their participants' sales. We have also probed the theoretical limits of this result by showing that, while in general it is possible for transparency to hinder collusion even in a completely stationary physical environment, this cannot occur if firms' profit functions satisfy a concavity condition that holds under linear Cournot and differentiated-product Bertrand competition. We therefore identify some settings where the standard view that transparency favors collusion is borne out, while also pointing out some little-recognized limitations of this view.

All of the results in this paper concern the comparison of information structures within a cartel: when is a cartel better-off when more or less information is exogenously available? A closely related question is that of how the desire to maintain privacy or transparency influences cartel behavior under a fixed information structure. From this perspective, we believe that our approach can offer a new explanation for the well-documented phenomenon of *price rigidity* in cartels, one which is quite different from existing approaches (Athey, Bagwell, and Sanchirico, 2004; Harrington and Chen, 2006). Consider an example similar to that in Sections 2.3 and 2.4: There are two firms, two markets, and two demand states, which are independent across markets and positively persistent across time. Prices are monitored perfectly. In a *flexible price equilibrium*, prices are tailored to current demand states: this has the advantage of allowing for higher profits in principle, but has the disadvantage of revealing the current demand state—and hence revealing information about future demand states—to one's competitor. In a *rigid price equilibrium*, prices are constant on-path. In this example, we have been able to show that, if the discount factor is intermediate and the gap between the low and high demand states is sufficiently large, then the best rigid price equilibrium yields higher profits than the best flexible price equilibrium. It seems quite plausible that the desire to maintain the privacy of one's home-market demand state is a rationale for rigid

pricing more generally. Developing this idea further is an interesting direction for future research.

More broadly, we hope to draw renewed attention to the role of information-sharing within cartels in supporting collusion. By assuming that cartel participants condition their behavior only on information that is common knowledge within the cartel, the existing theoretical literature on collusion has largely neglected the benefits that colluding firms can reap—at consumers' expense—from keeping their behavior private. Acknowledging the benefits as well as the costs of maintaining privacy in cartels may thus be a first step in improving our understanding of this aspect of antitrust economics.

Appendix A: Proof of Proposition 3

Under the first-best action plan, a firm's future profit when the previous demand state in its home market was low and high, respectively, is given by

$$\begin{aligned} V_L &= \frac{\tilde{s}_L^2}{4} + \delta\left[\phi V_L + (1-\phi) V_H\right] \\ V_H &= \frac{\tilde{s}_H^2}{4} + \delta\left[\phi V_H + (1-\phi) V_L\right]. \end{aligned}$$

Solving for V_L and V_H gives

$$\begin{aligned} V_L &= \frac{1}{(1-\delta)(1+\delta-2\delta\phi)}\left[(1-\delta\phi)\frac{\tilde{s}_L^2}{4} + \delta(1-\phi)\frac{\tilde{s}_H^2}{4}\right] \\ V_H &= \frac{1}{(1-\delta)(1+\delta-2\delta\phi)}\left[(1-\delta\phi)\frac{\tilde{s}_H^2}{4} + \delta(1-\phi)\frac{\tilde{s}_L^2}{4}\right]. \end{aligned}$$

As $\tilde{s}_L < \tilde{s}_H$, we have $V_L < V_H$. We also note that V_L and V_H are increasing in δ and go to infinity as $\delta \to 1$.

Suppose firms observe only industry demand. Then, as a deviator can be held to her minmax payoff of 0 (as we have seen), the first-best action plan is sequentially rational on path at $t > 0$ if and only if

$$V_L \geq \frac{\tilde{s}_L^2}{4} + (n-2)\max\left\{\frac{\tilde{s}_L(\tilde{s}_L - 2c)}{4}, \frac{\tilde{s}_H(\tilde{s}_H - 2c)}{8}\right\} + \frac{\tilde{s}_H(\tilde{s}_H - 2c)}{4}$$

and

$$V_H \geq \frac{\tilde{s}_H^2}{4} + (n-2)\max\left\{\frac{\tilde{s}_L(\tilde{s}_L - 2c)}{4}, \frac{\tilde{s}_H(\tilde{s}_H - 2c)}{8}\right\} + \frac{\tilde{s}_L(\tilde{s}_L - 2c)}{4},$$

or equivalently

$$\delta\left[\phi V_L + (1-\phi) V_H\right] \geq (n-2)\max\left\{\frac{\tilde{s}_L(\tilde{s}_L - 2c)}{4}, \frac{\tilde{s}_H(\tilde{s}_H - 2c)}{8}\right\} + \frac{\tilde{s}_H(\tilde{s}_H - 2c)}{4} \qquad (4)$$

and

$$\delta\left[\phi V_H + (1-\phi) V_L\right] \geq (n-2)\max\left\{\frac{\tilde{s}_L(\tilde{s}_L - 2c)}{4}, \frac{\tilde{s}_H(\tilde{s}_H - 2c)}{8}\right\} + \frac{\tilde{s}_L(\tilde{s}_L - 2c)}{4}.$$

As $\phi > 1/2$, $V_H > V_L$, and $\tilde{s}_H > \tilde{s}_L$, the former inequality implies the latter. Furthermore, noting that setting price $\bar{s}/2$ in all markets is the most tempting deviation at $t = 0$, sequential rationality holds at $t = 0$ if and only if

$$\delta \left[\frac{1}{2} V_L + \frac{1}{2} V_H \right] \geq (n-1) \frac{\bar{s} (\bar{s} - 2c)}{4}. \tag{5}$$

Hence, first-best industry profits are sustainable if and only if (4) and (5) hold. Finally, note that the left-hand sides of (4) and (5) are increasing in δ (and go to zero and infinity as $\delta \to 0$ and 1), and let δ^* be the cutoff value of δ such that one of (4) and (5) hold with equality while the other is satisfied.

Next, suppose firms observe all prices and sales. In this case, the first-best action plan is sustainable if and only if

$$\delta \left[\phi V_L + (1-\phi) V_H \right] \geq (n-2) \left(\frac{\tilde{s}_L (\tilde{s}_L - 2c)}{8} + \frac{\tilde{s}_H (\tilde{s}_H - 2c)}{8} \right) + \frac{\tilde{s}_H (\tilde{s}_H - 2c)}{4} \tag{6}$$

and (5) holds. However, as $\bar{s} = (\tilde{s}_L + \tilde{s}_H)/2$, Jensen's inequality implies that

$$\frac{\bar{s} (\bar{s} - 2c)}{4} < \frac{\tilde{s}_L (\tilde{s}_L - 2c)}{8} + \frac{\tilde{s}_H (\tilde{s}_H - 2c)}{8}.$$

Since $\phi > 1/2$, $V_H > V_L$, and $\tilde{s}_H > \bar{s}$, this means that (6) implies (5), so first-best industry profits are in fact sustainable if and only if (6) holds.

Finally, as $\tilde{s}_H > \tilde{s}_L$, the right-hand side of (6) is strictly greater than the right-hand side of (4). Hence, letting δ^{**} be the cutoff value of δ such that (6) holds with equality, we have $\delta^* < \delta^{**}$.

Appendix B: Entry Deterrence with a Non-Omniscient Mediator

We extend the construction of Section 4.2 to the case where the mediator does not directly observe the firms' actions or profits and must instead rely on the firms themselves to report these quantities. We do so by making two changes to the game from Section 4.2. First, we

augment the stage game by adding a mutual minmax Nash equilibrium action profile, so the payoff matrix is now

$$\begin{array}{cccc} & Out & In & PriceWar \\ Fight & 3,0 & -3,-3 & 0,0 \\ Accommodate & 4,0 & 0,3 & 0,0 \\ PriceWar & 0,0 & 0,0 & 0,0 \end{array} \tag{7}$$

Second, we assume that profits are stochastic conditional on actions (so payoff matrix (7) now represents expected rather than realized profits). Specifically, letting u_i denote firm i's expected profit function in (7), we assume that realized profits in period t are given by

$$\pi_{i,t}\left(a_t, \varepsilon_{i,t}\right) = u_i\left(a_t\right) + \varepsilon_{i,t},$$

where $\varepsilon_{i,t}$ is distributed uniformly on $[-1,1]$, independently across periods. We assume further that $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ are perfectly negatively correlated (so that $\varepsilon_{1,t} + \varepsilon_{2,t} = 0$ with probability 1) if the action profile is $(Fight, Out)$, but are independent otherwise. (This assumption implies that firm 1 loses information about firm 2's payoff following a deviation to *Accommodate*, which will be used to give firm 1 the required incentive to play $Fight$.)

The game in each period proceeds as follows:

1. The mediator makes a private action recommendation $r_{i,t} \in A_i$ to each firm i.
2. Firm i takes an action $a_{i,t} \in A_i$.
3. Firm i observes its realized profit $\pi_{i,t}$. The mediator observes nothing.
4. Firm i makes a private report of its profit $\hat{\pi}_{i,t}$ to the mediator.

Proposition 9 *In this augmented entry deterrence game with $\delta = 0.21$, there exists a sequential equilibrium with (non-omniscient) mediated private monitoring which yields a joint profit of* 3.5*.*

Proof. Consider the following strategy for the mediator:

- Say that a *deviation is detected in period t* if either of the following events occurs:

1. *Fight* is recommended to firm 1 and the firms' reports satisfy $\hat{\pi}_{1,t} + \hat{\pi}_{2,t} \neq 3$.
2. *Accommodate* is recommended to firm 1 and the firms' reports satisfy (at least one of) $\hat{\pi}_{1,t} < 2$, $\hat{\pi}_{2,t} > 1$, or $\hat{\pi}_{2,t} < -1$.

- If no deviation has been detected in any earlier period, issue recommendations as in the first two bullet points of the construction in the proof of Proposition 6.
- If a deviation has been detected in some earlier period, recommend $(PriceWar, PriceWar)$.

We show that it is an equilibrium for the firms to follow their action recommendations and report their profits truthfully.

Start by considering firms' incentives to misreport their profits. Note that the firms' reports matter only if no deviation has been detected so far: thus, in deciding what to report, the firms can condition on the event that no deviation has been detected. If firm 1 was recommended $Fight$, played $Fight$, and received profit $\pi_{1,t} \geq 2$, then she can infer that firm 2 played Out, so (given that firm 2 reports truthfully) with probability 1 only the truthful report $\hat{\pi}_{1,t} = \pi_{1,t}$ satisfies $\hat{\pi}_{1,t} + \hat{\pi}_{2,t} = 3$, and therefore firm 1 receives a positive continuation payoff from reporting truthfully and a zero continuation payoff from misreporting. If firm 1 was recommended $Fight$ but deviated or received profit $\pi_{1,t} < 2$, then she can infer that the action profile was not $(Fight, Out)$, so that (given that firm 2 reports truthfully) every possible report $\hat{\pi}_{1,t}$ will satisfy $\hat{\pi}_{1,t} + \hat{\pi}_{2,t} \neq 3$ with probability 1, so firm 1 may as well report truthfully. If firm 1 was recommended *Accommodate* and received profit $\pi_{1,t} \geq 2$, then she can infer that firm 2 played In, and a deviation will not be detected if she reports truthfully. Finally, if firm 1 was recommended *Accommodate* and received profit $\pi_{1,t} < 2$, then she can infer that firm 2 played Out, and a deviation will be detected whatever she reports (given that firm 2 reports truthfully).

Turning to firm 2, if firm 2 played Out, then if firm 1 played $Fight$ (and reports truthfully) then with probability 1 only the truthful report $\hat{\pi}_{2,t} = \pi_{2,t}$ satisfies $\hat{\pi}_{1,t} + \hat{\pi}_{2,t} = 3$, so truthful reporting is optimal; while if firm 1 played *Accommodate* then truthful reporting again does not lead to the detection of a deviation. If instead firm 2 played In, then (given that firm 1 reports truthfully) a deviation will be detected with probability 1 for either action of firm 1's and every report $\hat{\pi}_{2,t}$, so firm 2 may as well report truthfully. (This follows because if

firm 1 plays $Fight$ then $\hat{\pi}_{1,t} + \hat{\pi}_{2,t} \neq 3$ with probability 1, while if firm 1 plays *Accommodate* and reports truthfully then $\hat{\pi}_{1,t} < 2$.)

Finally, to see that following the mediator's action recommendations is optimal, note that the only tempting deviation is for firm 1 to deviate to *Accommodate* when recommended $Fight$. But, since $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ are independent when firm 1 plays *Accommodate*, the firms' reports will then satisfy $\hat{\pi}_{1,t} + \hat{\pi}_{2,t} \neq 0$ with probability 1. Since firm 1 was recommended $Fight$, this leads the mediator to recommend $(PriceWar, PriceWar)$ forever, which gives firm 1 a continuation payoff of 0. This implies that firm 1's incentive constraint is exactly the same as in the proof of Proposition 6, so this constraint is satisfied. ■

In fact, it is also possible to establish this result with the original payoff matrix (1), at the cost of a substantially more complicated equilibrium construction. The idea is to utilize a trick from our companion paper, Sugaya and Wolitzky (2016), whereby the mediator recommends every action profile with positive probability along the equilibrium path, so that a deviation by one firm is detectable by the mediator but not by the other firm. The other firm can then be induced to minmax the deviator without being made aware that a deviation has occurred. With this approach, it is possible to design a strategy profile where a deviant firm always receives its minmax payoff, despite the absence of a mutual minmax Nash equilibrium action profile. The details are available from the authors.

Appendix C: Linear Bertrand Competition

Consider n-player linear differentiated-product Bertrand competition, where firm i's demand at price vector p is given by

$$\max \left\{ a + b \sum_{j \neq i} p_j - p_i, 0 \right\}$$

for constants $a, b > 0$, so profits are given by

$$u_i(p) = \max \left\{ a + b \sum_{j \neq i} p_j - p_i, 0 \right\} p_i.$$

(Note that we allow negative prices but not negative quantities.) Restricting attention to equilibria with positive quantities along the equilibrium path, this game is concave (under

the weaker requirement that industry profits are concave) if and only if $b \leq 1/(n-1)$. To see this, first note that

$$V(p) = a\sum_i p_i + b\sum_i \sum_{j:i\neq j} p_i p_j - \sum p_i^2.$$

The Hessian matrix of $V(p)$ thus consists of -2's on the diagonal and $2b$'s off the diagonal, so this matrix is negative semi-definite if and only if

$$-\left(\sum_i x_i^2 - b\sum_i \sum_{j\neq i} x_i x_j\right) = -(1+b)\sum_i x_i^2 + b\left(\sum_i x_i\right)^2 \leq 0 \text{ for all } x \in \mathbb{R}^n.$$

This in turn holds if and only if $1+b \geq nb$, or $b \leq 1/(n-1)$. In addition, letting $\tilde{p}_{-i}$ denote a distribution over p_{-i}, we have

$$d(p_i, \tilde{p}_{-i}) = \frac{1}{4}\left(a + b\mathbb{E}^{\tilde{p}_{-i}}\left[\sum_{j\neq i} p_j\right]\right)^2 - \left(a + b\mathbb{E}^{\tilde{p}_{-i}}\left[\sum_{j\neq i} p_j\right] - p_i\right)p_i.$$

Hence, letting $g(p_{-i}) = \sum_{j\neq i} p_j$ and letting

$$f(p_i, x) = \frac{1}{4}(a+bx)^2 - (a + bx - p_i)p_i,$$

we have $d_i(p_i, \tilde{p}_{-i}) = f\left(p_i, \mathbb{E}^{\tilde{p}_{-i}}[g(p_{-i})]\right)$. The eigenvalues of the Hessian matrix of $f(p_i, x)$ are 0 and $2 + b^2/2$, so $f(p_i, x)$ is convex. In sum, the game is concave if and only if $b \leq 1/(n-1)$. Finally, unlike in the Cournot case, this game does not admit a mutual minmax Nash equilibrium. However, it follows from standard arguments (similar to Abreu, 1986) that it admits a subgame perfect equilibrium yielding payoff $\underline{u}$ whenever δ is sufficiently high. For example, this can be attained by a "stick-and-carrot" equilibrium path consisting of one period of large negative prices followed by an infinite stream of constant, positive prices. To complete the argument, we now verify that such a minmaxing stick-and-carrot equilibrium exists at a discount factor low enough such that first-best industry profits are unattainable—so that the conclusion of Proposition 8 is not trivial—whenever $b(n-1) < 0.933$.

To see this, note that the first-best price p^m is given by

$$\arg\max_p \left(a + b\left(n-1\right)p - p\right)p,$$

or

$$p^m = \frac{a}{2\left(1 - b\left(n-1\right)\right)}.$$

Corresponding firm profits u^m are given by

$$u^m = \left(a + b\left(n-1\right)p^m - p^m\right)p^m = \frac{a^2}{4\left(1 - b\left(n-1\right)\right)}.$$

Let δ_1 be lowest discount factor such that we can minmax a deviator with a stick-and-carrot equilibrium, where after a deviation all firms price at some level $\underline{p}$ for one period and then price at some level $\bar{p}$ forever. Let δ_2 be lowest discount factor such that price p^m is sustainable when deviators can be minmaxed. If $\delta_1 < \delta_2$, then Proposition 8 applies (and yields a non-trivial conclusion) whenever $\delta \in (\delta_1, \delta_2)$.

Proposition 10 *If* $b\left(n-1\right) < 0.933$ *then* $\delta_1 < \delta_2$.

Proof. In a stick-and-carrot equilibrium, the incentive compatibility constraint in the "carrot" state (pricing at $\bar{p}$) is

$$\left(a + b\left(n-1\right)\bar{p} - \bar{p}\right)\bar{p} \geq \left(1-\delta\right)\left(\frac{a + b\left(n-1\right)\bar{p}}{2}\right)^2 + \delta\left(0\right). \tag{8}$$

The incentive compatibility constraint in the "stick" state (pricing at $\underline{p}$) is

$$\left(1-\delta\right)\left(a + b\left(n-1\right)\underline{p} - \underline{p}\right)\underline{p} + \delta\left(a + b\left(n-1\right)\bar{p} - \bar{p}\right)\bar{p} \geq \left(\frac{a + b\left(n-1\right)\underline{p}}{2}\right)^2. \tag{9}$$

Also, if utility in the stick state equals the minmax payoff of 0, we must have

$$\left(1-\delta\right)\left(a + b\left(n-1\right)\underline{p} - \underline{p}\right)\underline{p} + \delta\left(a + b\left(n-1\right)\bar{p} - \bar{p}\right)\bar{p} = 0$$

and

$$a + b\left(n-1\right)\underline{p} = 0 \Leftrightarrow \underline{p} = \frac{a}{b\left(n-1\right)},$$

where the latter equation follows from (9).

Given this value of $\underline{p}$, we have

$$\left(a + b\left(n-1\right)\underline{p} - \underline{p}\right)\underline{p} = -\left(\frac{a}{b\left(n-1\right)}\right)^2,$$

so (9) becomes

$$\delta\left(a + b\left(n-1\right)\bar{p} - \bar{p}\right)\bar{p} \geq \left(1-\delta\right)\left(\frac{a}{b\left(n-1\right)}\right)^2. \tag{10}$$

This is a necessary and sufficient condition for the existence of a stick-and-carrot equilibrium with carrot price $\bar{p}$ that yields the minmax payoff in the stick state. Let δ_1 be the minimum discount factor for which there exists a price $\bar{p}$ that satisfies (8) and (10).

On the other hand, the condition for a constant price of p^m to be sustainable when deviators can be minmaxed is

$$\left(a + b\left(n-1\right)p^m - p^m\right)p^m \geq \left(1-\delta\right)\left(\frac{a + b\left(n-1\right)p^m}{2}\right)^2 + \delta\left(0\right). \tag{11}$$

Let δ_2 be the minimum discount factor that satisfies (11).

We wish to find a condition under which $\delta_1 < \delta_2$. First note that, if (10) holds with strict inequality at $\delta = \delta_2$ and $\bar{p} = p^m$, then $\delta_1 < \delta_2$. This follows because reducing $\bar{p}$ relaxes (8), so if (10) holds with strict inequality at $\bar{p} = p^m$ then there exists $\varepsilon > 0$ and $\delta < \delta_2$ such that both (8) and (10) are satisfied at discount factor δ when $\bar{p} = p^m - \varepsilon$.

Thus, it suffices to show that, when $\bar{p} = p^m$, the minimum discount factor at which (10) holds is less than the minimum discount factor at which (11) holds. Recalling the formula for u^m, rewrite (10) as

$$\delta\frac{a^2}{4\left(1 - b\left(n-1\right)\right)} \geq \left(1-\delta\right)\left(\frac{a}{b\left(n-1\right)}\right)^2,$$

or equivalently

$$\frac{1}{1-\delta} \geq 4\frac{1 - b\left(n-1\right)}{\left(b\left(n-1\right)\right)^2} + 1.$$

Similarly, noting that

$$\frac{a + b\left(n-1\right)p^m}{2} = \frac{a}{4}\left(\frac{2 - b\left(n-1\right)}{1 - b\left(n-1\right)}\right),$$

rewrite (11) as

$$\frac{1}{1-\delta} \geq \frac{1}{4}\frac{(2-b(n-1))^2}{1-b(n-1)}.$$

Hence, the minimum discount factor at which (10) holds is less than the minimum discount factor at which (11) holds if and only if

$$4\frac{1-b(n-1)}{(b(n-1))^2}+1 < \frac{1}{4}\frac{(2-b(n-1))^2}{1-b(n-1)}.$$

Letting $z = b(n-1)$, this is equivalent to

$$16 < \frac{(2-z)^2 z^2}{(1-z)(1-z+z^2)}.$$

This holds for all $z < 0.933$. ■

References

[1] Abreu, D. (1986), "Extremal Equilibria of Oligopolistic Supergames," *Journal of Economic Theory*, 39, 191-225.

[2] Abreu, D., P. Milgrom, and D. Pearce (1991), "Information and Timing in Repeated Partnerships," *Econometrica*, 1713-1733.

[3] Abreu, D., D. Pearce, and E. Stacchetti (1986), "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory*, 39, 251-269.

[4] Albaek, S., P. Møllgaard, and P.B. Overgaard (1997), "Government-Assisted Oligopoly Coordination? A Concrete Case," *Journal of Industrial Economics*, 45, 429-443.

[5] Aoyagi M. (2002), "Collusion in Dynamic Bertrand Oligopoly with Correlated Private Signals and Communication," *Journal of Economic Theory*, 102, 229-248.

[6] Aoyagi, M. (2007), "Efficient Collusion in Repeated Auctions with Communication," *Journal of Economic Theory*, 134, 61-92.

[7] Areeda, P. and L. Kaplow (1997), *Antitrust Analysis, Fifth Edition*, New York: Aspen.

[8] Athey, S. and K. Bagwell (2001), "Optimal Collusion with Private Information," *RAND Journal of Economics*, 32, 428-465.

[9] Athey, S., K. Bagwell, and C. Sanchirico (2004), "Collusion and Price Rigidity," *Review of Economic Studies*, 71, 317-349.

[10] Awaya, Y. and V. Krishna (2015), "On Communication and Collusion," *American Economic Review*, forthcoming.

[11] Bagwell, K. and R.W. Staiger (1997), "Collusion over the Business Cycle," *RAND Journal of Economics*, 28, 82-106.

[12] Bernheim, B.D. and E. Madsen (2016), "Price Cutting and Business Stealing in Imperfect Cartels," *mimeo*.

[13] Blackwell, D. (1951), "Comparison of Experiments," *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, 1, 93-102.

[14] Carlton, D.W. and J.M. Perloff (1995), *Modern Industrial Organization, Fourth Edition*, Boston: Pearson

[15] Chang, M.-H. (1991), "The Effects of Product Differentiation of Collusive Pricing," *International Journal of Industrial Organization*, 3, 453-470.

[16] Compte, O. (1998), "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 66, 597-626.

[17] Compte, O., F. Jenny, and P. Rey (2002), "Capacity Constraints, Mergers and Collusion," *European Economic Review*, 46, 1-29.

[18] Forges, F. (1986), "An Approach to Communication Equilibria," *Econometrica*, 54, 1375-1385.

[19] Friedman, J.W. (1971), "A Non-Cooperative Equilibrium for Supergames," *Review of Economic Studies*, 38, 1-12.

[20] Genesove, D. and W.P. Mullin (2001), "Rules, Communication, and Collusion: Narrative Evidence from the Sugar Institute Case," *American Economic Review*, 91, 379-398.

[21] Goldlücke, S. and S. Kranz (2012), "Infinitely Repeated Games with Public Monitoring and Monetary Transfers," *Journal of Economic Theory*, 147, 1191-1221.

[22] Green, E.J. and R.H. Porter (1984), "Noncooperative Collusion under Imperfect Price Information," *Econometrica*, 52, 87-100.

[23] Häckner, J. (1996), "Optimal Symmetric Punishments in a Bertrand Differentiated Products Duopoly," *International Journal of Industrial Organization*, 14, 611-630.

[24] Haltiwanger, J. and J.E. Harrington (1991), "The Impact of Cyclical Demand Movements on Collusive Behavior," *RAND Journal of Economics*, 22, 89-106.

[25] Harrington, J.E. (2006), "How do Cartels Operate?" *Foundations and Trends in Microeconomics*, 2, 1-105.

[26] Harrington, J.E. and J. Chen (2006), "Cartel Pricing Dynamics with Cost Variability and Endogenous Buyer Detection," *International Journal of Industrial Organization*, 24, 1185-1212.

[27] Harrington, J.E. and A. Skrzypacz (2011), "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices," *American Economic Review*, 101, 2425-2449.

[28] Hirshleifer, J. (1971), "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 61, 561-574.

[29] Hochman, G. and E. Segev (2010), "Managed Trade with Imperfect Information," *International Economic Review*, 51, 187-211.

[30] Hörner, J. and A. Jamison (2007), "Collusion with (Almost) No Information," *RAND Journal of Economics*, 38, 604-822.

[31] Kandori, M. (1991a), "Cooperation in Finitely Repeated Games with Imperfect Private Information," *mimeo*.

[32] Kandori, M. (1991b), "Correlated Demand Shocks and Price Wars During Booms," *Review of Economic Studies*, 58, 171-180.

[33] Kandori, M. (1992), "The Use of Information in Repeated Games with Imperfect Monitoring," *Review of Economic Studies*, 59, 581-593.

[34] Kandori, M. (2002), "Introduction to Repeated Games with Private Monitoring," *Journal of Economic Theory*, 102, 1-15.

[35] Kandori, M. and H. Matsushima (1998), "Private Observation, Communication and Collusion," *Econometrica*, 66, 627-652.

[36] Lambson, V.E. (1987), "Optimal Penal Codes in Price-Setting Supergames with Capacity Constraints," *Review of Economic Studies*, 54, 385-397.

[37] Levenstein, M.C. and V.Y. Suslow (2006), "What Determines Cartel Success?" *Journal of Economic Literature*, 44, 43-95.

[38] Mailath, G.J., S.A. Matthews, and T. Sekiguchi (2002), "Private Strategies in Finitely Repeated Games with Imperfect Public Monitoring," *Contributions in Theoretical Economics*, 2.

[39] Marshall, R.C. and L.M. Marx (2012), *The Economics of Collusion*, Cambridge: MIT Press.

[40] Miyahara, Y. and T. Sekiguchi (2013), "Finitely Repeated Games with Monitoring Options," *Journal of Economic Theory*, 148, 1929-1952.

[41] Myerson, R.B. (1986), "Multistage Games with Communication," *Econometrica*, 54, 323-358.

[42] Pesendorfer, M. (2000), "A Study of Collusion in First-Price Auctions," *Review of Economic Studies*, 67, 381-411.

[43] Rahman, D. (2014), "The Power of Communication," *American Economic Review*, 104, 3737-3751.

[44] Ross, T. (1992), "Cartel Stability and Product Differentiation," *International Journal of Industrial Organization*, 10, 1-13.

[45] Rotemberg, J.J. and G. Saloner (1986), "A Supergame-Theoretic Model of Price Wars During Booms," *American Economic Review*, 76, 390-407.

[46] Sekiguchi, T. (2002), "Existence of Nontrivial Equilibria in Repeated Games with Imperfect Private Monitoring," *Games and Economic Behavior*, 40, 299-321.

[47] Sorin, S. (1986), "On Repeated Games with Complete Information," *Mathematics of Operations Research*, 11, 147-160.

[48] Stigler, G.J. (1964), "A Theory of Oligopoly," *Journal of Political Economy*, 72, 44-61.

[49] Sugaya, T. and A. Wolitzky (2016), "Bounding Equilibrium Payoffs in Repeated Games with Private Monitoring," *Theoretical Economics*, forthcoming.

[50] Wernerfelt, B. (1989), "Tacit Collusion in Differentiated Cournot Games," *Economics Letters*, 29, 303-306.

[51] Whinston, M.D. (2006), *Lectures on Antitrust Economics*, Cambridge: MIT Press.