

MIT Open Access Articles

Maintaining Privacy in Cartels

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Sugaya, Takuo and Alexander Wolitzky. "Maintaining Privacy in Cartels." *Journal of Political Economy* 126, 6 (December 2018): 2569-2607 © 2018 University of Chicago.

As Published: 10.1086/699975

Publisher: University of Chicago Press

Persistent URL: <https://hdl.handle.net/1721.1/122787>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Maintaining Privacy in Cartels

Takuo Sugaya

Stanford University

Alexander Wolitzky

Massachusetts Institute of Technology

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 worked to preserve the privacy of their participants' actions and outcomes. Toward explaining this behavior, we show that cartels can sometimes sustain higher profits when actions and outcomes are observed only privately, because better information can hinder collusion by helping firms devise more profitable deviations from the collusive agreement. We provide conditions under which maintaining privacy is optimal for cartels that follow a market-segmentation strategy.

I. 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, 40)

We thank John Asker, Leslie Marx, and Tom Wiseman for insightful discussions of this paper. For helpful comments, we thank Daron Acemoglu, Nikhil Agarwal, Yu Awaya, Kyle Bagwell, Glenn Ellison, Joe Harrington, Ali Hortaçsu, Louis Kaplow, Vijay Krishna, Kevin Li, Markus Möbius, Juan Ortner, Andy Skrzypacz, Juuso Toikka, Glen Weyl, and Mike Whinston; participants in the 2016 Warwick Economic Theory Workshop, the 2016 Searle Antitrust Conference, the 2016 Columbia/Duke/MIT/Northwestern IO Theory Conference, the Winter 2017 National Bureau of Economic Research Industrial Organization Meetings, and seminars at Microsoft Research, Massachusetts Institute of Technology, Northwestern,

Electronically published October 31, 2018

[*Journal of Political Economy*, 2018, vol. 126, no. 6]

© 2018 by The University of Chicago. All rights reserved. 0022-3808/2018/12606-0009\$10.00

puts it in his monograph on antitrust economics, “Lesser observability, including more noisy signals of price cuts, makes sustaining a given supra-competitive 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, 136) and antitrust law (“[To sustain collusion,] firms must be able to observe and compare each others’ prices”; Areeda and Kaplow 1997, 254) and is a prominent part of the US Department of Justice/Federal Trade Commission horizontal-merger guidelines (“A market typically is more vulnerable to coordinated conduct if each competitively important firm’s significant competitive initiatives can be promptly and confidently observed by that firm’s rivals”; USDoJ/FTC 2010, 26). The theory has also been successfully applied in several well-known empirical studies, such as Albæk, Møllgaard, and Overgaard’s (1997) work on the Danish ready-mixed concrete industry and Genesove and Mullin’s (2001) study of the prewar US sugar industry.

There are also, however, various pieces of evidence that suggest that the conventional wisdom may not tell the whole story. Most strikingly, several recent cartels uncovered by the European Commission (EC) 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 for copper plumbing tubes, and a consulting firm for low-density polyethylene), which then returned only aggregate statistics to the firms.¹ This behavior is a puzzle for the view that transparency facilitates collusion; as Harrington (2006, 54) writes, “It is unclear why firms sought to maintain privacy of their market share and to what extent effective enforcement could be achieved without market shares being commonly known among the cartel members.”

There are also theoretical reasons why cartels might strive to maintain privacy rather than transparency. It is a familiar idea in economics that giving agents too much information can hurt their incentives to cooper-

New York University, Michigan State, Oxford, Stanford, Tokyo, and University of Pennsylvania; and the anonymous referees. A. Wolitzky acknowledges financial support from the National Science Foundation.

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).

ate 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. We argue that there are (at least) three major effects of increased transparency on the sustainability of collusion:

- Effect M (monitoring). More information makes it easier to detect deviations from the collusive agreement.
- Effect C (coordination). More information helps the cartel tailor collusive prices to current market conditions.
- Effect D (deviations). More information lets individual firms tailor deviations to current market conditions.³

As we discuss at length below, effects M and C are well understood and suggest that transparency facilitates collusion. Effect D is more novel—especially when it refers to information about firms’ actions, rather than information about payoff-relevant parameters—and it goes the opposite way. The main contributions of this paper are to (1) identify a natural setting where effect D is dominant, so that transparency hinders collusion, (2) argue that this is a plausible explanation for several real-world cartels’ efforts to maintain the privacy of their participants’ actions, and (3) provide some preliminary analysis of the trade-offs among the three effects, along with comparative statics on how the optimal degree of transparency varies with other aspects of the economic environment.

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 We thank a referee for suggesting this decomposition of the role of information. The decomposition is not exhaustive. For example, in some settings additional information will also help the cartel tailor off-path punishments to market conditions, providing another force in favor of transparency.

We consider a setting where the global market is segmented by geographic or product characteristics, each firm has a cost advantage in its home market, and cost and demand conditions are stochastic. In this environment, the joint 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).⁴ 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 United States, 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 (2006, 34) 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."⁵

Our main result provides conditions under which 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, it is reasonable to assume that entry into one's home market is detectable and that each firm does not need to know the demand state in the "foreign" markets to price optimally in its home market. These assumptions—discussed further below—effectively rule out effects M and C above. 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 (effect D). 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

⁴ Interestingly, Stigler (1964) noted that market segmentation could be an effective form of collusion but thought it would be easily detected and forestalled by antitrust authorities. Market segmentation is, of course, related to the large literature on multimarket contact, following Bernheim and Whinston (1990). Two strands of this literature are particularly relevant to our study. First, Matsushima (2001), Kobayashi and Ohta (2012), and Sekiguchi (2015) study multimarket contact with imperfect monitoring. Second, Belleflamme and Bloch (2008), Byford and Gans (2014), and Bhattacharjea and Sinha (2015) explicitly study market segmentation (with perfect monitoring).

⁵ 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 Sec. V for further discussion. Many other examples are available. For instance, Pesendorfer (2000) reports that the Texas school milk cartel operated by dividing the market on the basis of cost advantages and refraining from entry in others' home areas, while occasionally swapping contracts among firms in response to changing costs; and Asker (2010) documents explicit and persistent customer allocation in the parcel tanker shipping industry.

fails badly in stochastic environments where the cartel tries to segment the market.

We provide 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 EC 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. We also find (weak) evidence of a positive statistical correlation between these features.

Finally, we also present several theoretical examples aimed at clarifying the boundaries of our core model and mechanism.

We first give parameterized examples where effects M and C are present in addition to effect D, so there are both benefits and costs of information. In these examples, we show that the cartel-optimal level of information is increasing in the discount factor and decreasing in the persistence of market demand, thus providing testable comparative statics predictions.

Perhaps most surprisingly, we also show that effect D can be present (and that maintaining privacy can thus be required for sustaining collusion) even when the physical environment is completely stationary, so that the model is a standard repeated game 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 perfect 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 be essential for supporting collusion in a special case of our model with a stationary physical environment. The intuition is that—despite the stationarity of the physical environment—a collusive equilibrium must sometimes have a nonstationary 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.

The reader may wonder how, 50 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.⁶ However, in the more realistic case where firms receive private signals, this intuition breaks 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 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 information (Athey and Bagwell 2001; Aoyagi 2002; Harrington and Skrzypacz 2011; Rahman 2014; Chan and Zhang 2015; Spector 2015; Awaya and Krishna 2016). In particular, Rahman, Spector, 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.⁷

A recent paper by Kloosterman (2015) also makes the point that more information can make cooperation harder in a Markov game. However, Kloosterman assumes perfect monitoring of actions and perfect information about the current physical environment and examines the impact of more precise public information regarding the next period’s physical environment. Thus, while the negative effect of information we emphasize is that better information (including better information about actions alone) lets players devise more profitable deviations, Kloosterman’s point is that better information about tomorrow’s physical environment can prompt deviations by delivering bad news about equilibrium continuation payoffs. The former effect is absent in Kloosterman’s setting (because of his assumption of perfect information), and the latter effect is absent in

⁶ 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. Fuchs (2007) present a related result in a repeated principal-agent model.

⁷ Athey and Bagwell (2001) do present a numerical example where firms benefit from limiting communication about payoff-relevant state variables.

ours (under the assumptions of our main result, theorem 1). The papers are thus complementary.

Finally, a companion paper (Sugaya and Wolitzky 2017) contains an example where players in an infinitely repeated game benefit from imperfections in the monitoring technology.⁸ The example is slightly related to example 4 (Sec. IV.D) here. The main result of the companion paper is a sufficient condition for transparency (i.e., perfect monitoring) to be the optimal information structure for sustaining cooperation. A second paper (Sugaya and Wolitzky 2018a) provides a different sufficient condition, which in particular implies that transparency is optimal in repeated Cournot or differentiated-product Bertrand competition with linear demand curves. In contrast, the main result of the current paper gives sufficient conditions for privacy to be optimal.

The remainder of the paper is organized as follows. Section II develops the model and reprises the standard result that transparency facilitates collusion with imperfect public monitoring. Section III presents our main theorem: under some assumptions, transparency hinders collusion with private monitoring. Section IV contains examples that complement the theorem in various ways and in particular address the trade-off among the three effects of improved information discussed above. Finally, Section V discusses several real-world cartels through the lens of our model, and Section VI concludes. The proof of the main theorem is contained in appendix A; all other omitted proofs may be found in appendix C (apps. B and C are available online).

II. Model

We consider a fairly general model of multimarket price competition with homogeneous goods and stochastic costs and demand.

A. Formal Model

1. Physical Environment and Payoffs

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 constitute the demand side of the market. In every period $t = 0, 1, 2, \dots$, each firm i can produce in market j at constant marginal cost $c_i^j \geq 0$, where the vector of cost states $c = (c_i^j) \in (C_i^j) = C$ can change over time as described below. We assume that $c_i^j \leq$

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

c_i^j 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 .⁹

In every period t , firm i chooses a price vector $(p_i^j)_{j=1}^n \in (\mathbb{R}_+ \cup \{\infty\})^n$, where p_i^j is firm i 's price in market j . As we see below, 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 $D(\underline{p}^j, s^j)$, 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 $s \in S = (S^j)$ denote a vector of market demand states. Assume that the function $D(p^j, s^j)$ is continuous, nonnegative, and strictly increasing in s^j , with $D(p^j = c_j^j, s^j) > 0$ for all $(c_j^j, s^j) \in C_j^j \times S^j$, $D(p^j, s^j)$ p^j bounded, and $D(\infty, s^j) = 0$.¹⁰ The lowest-price firms in market j supply all $D(\underline{p}^j, s^j)$ units at price \underline{p}^j , with the market allocated to the home firm in case of a tie (or allocated arbitrarily among the lowest-price firms if the home firm does not have the lowest price).¹¹ Denote the vector of sales at price vector p and demand state vector s $q(p, s) = (q_i^j(p^j, s^j))$. Finally, assume that the sets of cost and demand states C and S are finite and that the vector $(c_t, s_t) \in C \times S$ follows a Markov process, with Markov transition function $M: C \times S \rightarrow \Delta(C \times S)$. The prior distribution over period 0 cost and demand states is given by $\varrho \in \Delta(C \times S)$.

Firm i 's period t profit, given a (price, cost, demand) vector (p_t, c_t, s_t) , is thus given by

$$u_{i,t} := \sum_{j=1}^n q_i^j(p_t^j, s_t^j) (p_{i,t}^j - c_{i,t}^j).$$

Each firm maximizes its discounted expected profit $\sum_{t=0}^{\infty} \delta^t E[u_{i,t}]$, where $\delta < 1$ is the common discount factor.

2. Information Structure

At the beginning of each period t , firms observe their own period $t - 1$ prices, period $t - 1$ sales, and period t costs and can also receive signals about the entire vector of period $t - 1$ prices, period $t - 1$ demand states, and period t costs.¹² Specifically, there is a finite set of signals $Z = (Z_i)$ and a family of conditional probability distributions (an “information structure”) on Z , $\pi(z|p, c', s)$, such that signal z is realized with

⁹ For a recent model of dynamic price competition with a similar distinction between home and foreign markets, see Bernheim and Madsen (2017).

¹⁰ We follow the convention $0 \cdot \infty = 0$.

¹¹ This tie-breaking rule would emerge endogenously in Bertrand competition with heterogeneous costs if consumers could choose from whom to purchase.

¹² Consistent with this notation, players start the game knowing their own period 0 costs and receive signals about others' period 0 costs only.

probability $\pi(z|p, c', s)$ when p is the vector of prices and (c', s) is the vector of current costs and last period's demand states.¹³ When signal z is realized, firm i observes only its i th component, z_i . Let π_i denote the marginal distribution of π over Z_i .

Examples of information structures include the cases where firms observe the vector of prices p , cases where firms observe the vector of sales $q(p, s)$, and cases where firms observe only total industry sales $\sum_{i,j} q_i^j(p^j, s^j)$. We consider these and other cases in the context of parameterized examples in Section IV.

3. Solution Concept

Unless otherwise specified, the solution concept is weak perfect Bayesian equilibrium (hereafter “equilibrium”).

B. A Benchmark Result: Transparency Facilitates Collusion with Observable States and Public Monitoring

In terms of the effects of information discussed in Section I, our main result gives conditions under which effects M and C are inactive and effect D implies that transparency hinders collusion. Before presenting this result, we first give a version of the standard result, where effects C and D are inactive and effect M implies that transparency helps collusion.

To shut down effects C and D, assume that the physical state is perfectly observed (formally, if $\pi_i(z_i|p, c, s) > 0$, then $\pi_i(z_i|p, \hat{c}, \hat{s}) = 0$ for all $(\hat{c}, \hat{s}) \neq (c, s)$) and that monitoring is public ($\pi(z|p, c, s) = 0$ whenever $z_i \neq z'_i$ for some i, i'). To capture the standard intuition for why transparency helps collusion, it is also necessary to restrict attention to *perfect public equilibria*, which are perfect Bayesian equilibria in which play in period t is conditioned only on the publicly available history of signals $(z_\tau)_{\tau=0}^t$, and to assume that the players have access to a public randomization device.¹⁴ Finally, recall that an information structure π' is (*Blackwell*) *more informative* than π for every firm i (denoted $\pi' \geq \pi$) if for each i there exists a function $f_i: Z_i \times Z_i \rightarrow [0, 1]$, such that $\sum_{z_i \in Z_i} f_i(z_i, z'_i) = 1$ for all $z'_i \in Z_i$ and that, for all $z_i \in Z_i, p, c, s$,

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

We say that π' is *strictly more informative* than π (denoted $\pi' > \pi$) if, in addition, $f_i(z_i, z'_i) > 0$ for all $z_i, z'_i \in Z_i$ for every firm i .

¹³ Note that this formulation allows signals of sales, as sales are determined by prices and demand states.

¹⁴ For details, see Abreu, Pearce, and Stacchetti (1990), Kandori (1992), and Fudenberg, Levine, and Maskin (1994); or see the textbook treatment in Mailath and Samuelson (2006).

PROPOSITION 0. With observable states, public monitoring, and the availability of public randomization, making the information structure more informative weakly expands the perfect public equilibrium payoff set. In particular, maximum industry profits are nondecreasing in the informativeness of the information structure.

Proof. This is a straightforward extension of proposition 1 of Kandori (1992) from repeated games to Markov games. It also follows immediately from corollary 1 of Kim (2016). QED

Proposition 0 is one version of the standard result that transparency helps collusion. It relies on the assumptions that all payoff-relevant states of the world are observable and that firms condition their behavior only on publicly available information. As we see below, when these special assumptions are replaced by other (equally special) assumptions that turn off effect M rather than effect D, we find the opposite result, namely, that transparency hinders collusion.

III. Main Result: Conditions for Transparency to Hinder Collusion with Imperfectly Observed States and Private Monitoring

This section contains our main result: under conditions on the information structures that have the effect of shutting down effects M and C, 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. While the required conditions are quite special, we believe that they are reasonable in many contexts where a cartel sells homogeneous goods in a segmented market.

First and foremost, we assume that a firm can perfectly distinguish between situations in which another firm enters its home market and situations in which this does not occur.

ASSUMPTION 1. For every market j and all 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 $i \neq j$, the following holds: for every $z_j \in Z_j$, if $\pi_j(z_j|p, c, s) > 0$ for some $(c, s) \in C \times S$, then $\pi_j(z_j|\hat{p}, c, s) = 0$ for all $(c, s) \in C \times S$.

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 that it has not tried to sell in market j . 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 into one's home market. For example, Harrington (2006, 46) writes that "an attraction to a customer allocation 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.” Marshall and Marx (2012, 131) add that “with a geographic allocation, if each producer is in a separate country and information about cross-border trade is readily available . . . then monitoring can be straightforward.” As a practical example, in Section V we describe how such monitoring worked in the European copper plumbing tubes cartel.

In any case, when the cartel tries to segment the market, assumption 1 is the key assumption that shuts down effect M: under assumption 1, no further information about one’s competitors’ prices, sales, or profits is necessary to detect deviations from market segmentation.

We also assume that signals have a product structure, which allows every signal z_i to be decomposed into signals $x_{j,i}^k$ 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’s costs. This means that signals of some prices are not directly informative about either other prices or the demand and cost states—although, of course, in equilibrium firms will draw inferences about demand and costs based on signals of prices.

ASSUMPTION 2. There exist finite sets $X_j^k = (X_{j,i}^k)$ and $Y = (Y_i)$ and families of conditional probability distributions $\pi^{X_j^k}(x_j^k | p_j^k)$ and $\pi^Y(y | p, c, s)$ such that

1. $Z = (\prod_{j,k} X_j^k) \times Y$,
2. $\pi((x, y) | p, c, s) = (\prod_{j,k} \pi^{X_j^k}(x_j^k | p_j^k)) \pi^Y(y | p, c, s)$, and
3. $\pi^Y(y | p, c, s) = \pi^Y(y | \hat{p}, c, s)$ whenever $q(p, s) = q(\hat{p}, s)$.

Assumption 2 plays a smaller and more technical role in the analysis, which is discussed below.

We further assume that cost and demand transitions at the level of an individual firm or market depend only on that firm’s or market’s current costs and demand.¹⁵ The economic content of this 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 home market. This assumption shuts down effect C.

ASSUMPTION 3. For every firm i , there is a function $M_i: C_i \times S^i \rightarrow \Delta(C_i \times S^i)$ such that $M_i(c_i^t, s^t) = M(c, s)|_{C_i \times S^i}$ for all $(c, s) \in C \times S$, where $M(c, s)|_{C_i \times S^i}$ denotes the projection of $M(c, s)$ onto $C_i \times S^i$.

If firm i were a monopoly in market i and observed the period $t - 1$ vector of demand states s_{t-1} and the period t cost vector c_t , it would set price p_t^i in period t to maximize $\mathbb{E}[D(p_t^i, s_t^i)(p_t^i - c_{i,t}^i) | c_t, s_{t-1}]$, which equals

¹⁵ Note that this assumption does not imply that state transitions are independent across firms or markets.

$\mathbb{E}[D(p_i^i, s_i^i)(p_i^i - c_{i,t}^i) | c_{i,t}^i, s_{i-1}^i]$ by assumption 3. Let $p_i^m(c_{i,t}^i, s_{i-1}^i)$ be a solution to this problem. Let $p_i^m(c_{i,0}^i, \emptyset)$ be a maximizer of $\mathbb{E}[D(p_i^i, s_0^i)(p_i^i - c_{i,0}^i) | c_{i,0}^i]$.

Because of 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.

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

(In terms of social welfare, note that the home-market principle entails productive efficiency but of course also involves the usual monopoly quantity distortion in each market.)

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

ASSUMPTION 4. For all i , $z_i \in Z_i$, and $c \neq \hat{c}$, if $\pi_i(z_i | p, c, s) > 0$ for some (p, s) , then $\pi_i(z_i | p, \hat{c}, s) = 0$ for all (p, s) .

Finally, we also give conditions under which collusion is strictly more difficult to sustain under a strictly more informative information structure. This result requires two additional technical assumptions.

ASSUMPTION 5. The following full-support conditions hold: $\varrho(c, s) > 0$ for all $(c, s) \in C \times S$; $M(c, s | \hat{c}, \hat{s}) > 0$ for all $(c, s), (\hat{c}, \hat{s}) \in C \times S$; and $\pi(z | p, c, s) > 0$ for all $z \in Z$, price vectors p , and $(c, s) \in C \times S$.

ASSUMPTION 6. The prior belief $\varrho \in \Delta(C \times S)$ lies in the interior of the convex hull of the set of beliefs over $C \times S$ that arises in equilibrium when firms follow the home-market principle.¹⁶

The following is our main result.

THEOREM 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 informative than π , then $\delta^*(\pi') \geq \delta^*(\pi)$. In this sense, a more informative information structure hinders collusion. In addition, under assumptions 1–6, if π' is strictly more informative than π , then $\delta^*(\pi') > \delta^*(\pi)$.

¹⁶ This assumption is stated more formally in app. A.

The intuition for theorem 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 and punish deviations, as a firm can always detect entry into its home market (by assumption 1) and can always punish a deviator by pricing at its cost (by assumption 4). 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 that are better tailored to foreign-market conditions—and hence are more profitable. Providing this information thus increases the discount factor required for sustaining collusion.¹⁷

The proof of theorem 1 formalizes this intuition by characterizing the set of beliefs about foreign-market demand states that can arise in equilibrium when firms follow the home-market principle. We apply a fixed-point characterization of equilibrium beliefs due to Phelan and Skrzypacz (2012) to show that this set of beliefs expands with a more informative information structure, and (under assumptions 5 and 6) strictly expands when the information structure is strictly more informative.¹⁸ As a firm's maximum deviation gain is convex in its beliefs, this argument shows that the maximum deviation gain is increasing in the informativeness of the information structure, thus formalizing effect D. Finally, since the possible benefits of improved information are shut down, this increase in the maximum deviation gain implies that the firms must be more patient to sustain collusion.

Of course, in reality information about competitors' prices, costs, and sales does often provide additional information about deviations (effect M) and own-market conditions (effect C). Indeed, we analyze the trade-off between these effects and the more novel effect D in parameterized examples in Sections IV.B and IV.C. Thus, the point of theorem 1 is simply to highlight an opposing force favoring privacy that, in the context of a particular market application, would have to be weighed against the well-known advantages of transparency.

We conclude this section with a comment on two of the assumptions underlying theorem 1: the assumption that costs are observable (assumption 4) and the assumption that firms are willing to price below cost out-

¹⁷ 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 supporting less precise monitoring of other variables.

¹⁸ Phelan and Skrzypacz develop their results in the context of repeated games. We show how to extend their results to Markov games in the course of the proof of theorem 1.

side their home markets to punish deviators, so long as they do not expect to make any sales at these unprofitable prices. Interestingly, theorem 1 continues to hold if we instead make the opposite assumptions, namely, that costs are completely unobservable (and in particular cannot be inferred from information about prices or demand) and that firms never post unprofitable prices.

To this end, we impose the following conditions.

ASSUMPTION 7.

1. There exist functions $M^C : C \rightarrow \Delta(C)$ and $M^S : S \rightarrow \Delta(S)$ such that $M(c, s) = M^C(c)M^S(s)$ for all $(c, s) \in C \times S$.
2. $\pi^{X^i}(x_j^k | p_j^k) = \pi^{X^i}(x_j^k | \hat{p}_j^k)$ 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 , and \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 , and price p_i^j played with positive probability at any information set in period t .

In a cautious equilibrium, the harshest punishment for firm j involves every firm i pricing at $c_{i,t}^j$ in market j . Assumption 7 ensures that firm j does not obtain any information (other than its own costs) about how the severity of this punishment evolves over the course of the game.

PROPOSITION 1. Under assumptions 1–3 and 7, 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 informative than π , then $\tilde{\delta}(\pi') \geq \tilde{\delta}(\pi)$.

IV. Illustrative Examples

The remainder of the analysis consists of four examples that illustrate the applicability and the boundaries of theorem 1. The main conclusions of the examples may be roughly summarized as follows.

Example 1. When there is no uncertainty regarding aggregate industry demand, letting firms observe their competitors' individual prices and sales in addition to industry sales can strictly hinder collusion. (This is just a simple illustration of theorem 1 in a context where the information structure matches that in the industrial cartels discussed in the introduction and in Sec. V.)

Example 2. If entry into one's home market is not perfectly detectable—so that both effects M and D are present—then the trade-off between improved monitoring (effect M) and firms' ability to devise more profitable deviations (ef-

fect D) can make an intermediate level of transparency optimal for sustaining collusion. In this setting, the optimal level of transparency is decreasing in the persistence of demand.

Example 3. If information about foreign-market conditions is relevant for optimal monopoly pricing in the home market—so that now effects C and D are present—then an intermediate level of transparency can again be optimal. In this case, the optimal level of transparency is increasing in the discount factor.

Example 4. Perhaps most surprisingly, transparency can strictly hinder collusion even if the physical environment is completely stationary.

Finally, we also consider a slight extension of the model in appendix B: while our main model and results concern the case of homogeneous products, we show there that the logic of theorem 1 also applies in some settings with differentiated products.

A. *Example 1: Observing Firm-Level Sales Can Strictly Hinder Collusion*

Add the following assumptions to those imposed in Section III:

- The number of firms and markets n is even and at least 4.
- Costs are constant, with $c_i^i = 0$ and $c_i^j = c > 0$ for all $i \neq j$.
- Market demand curves are linear: $D(p^j, s^j) = s^j - p^j$.
- 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 $0 < \phi < 1$. 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.
- Firms observe only their own prices and sales, as well as total industry sales $\sum_j D(p^j, s^j)$.

While the assumption that firms perfectly observe total industry sales is consistent with the information-sharing practices of many of the cartels discussed in the introduction and Section V, it is not consistent with the full-support assumption (assumption 5) used to show that transparency can strictly hinder collusion. In addition, the assumption that de-

mand in adjacent markets is perfectly negatively correlated—so that aggregate industry demand is held constant—also violates the full-support assumption. The main contribution of this example is to show that these assumptions offset each other: in the absence of aggregate demand uncertainty, observing individual-level sales in addition to industry sales can strictly hinder collusion, even if industry sales are already perfectly observable.¹⁹

Let $\tilde{s}_L = \phi s_{L_1} + (1 - \phi) s_{H_1}$, $\tilde{s}_H = (1 - \phi) s_{L_2} + \phi s_{H_2}$, and $\bar{s} = (s_{H_1} + s_{L_2})/2$. Monopoly prices are given by $p_i^m(s_{L_1}) = \tilde{s}_L/2$, $p_i^m(s_{H_1}) = \tilde{s}_H/2$, and $p_i^m(\emptyset) = \bar{s}/2$. Let $\tilde{s}_{\min} := \min\{\tilde{s}_L, \tilde{s}_H\}$ and $\tilde{s}_{\max} := \max\{\tilde{s}_L, \tilde{s}_H\}$. Finally, assume that $\min\{s_{L_1}, s_{H_1}\} > \tilde{s}_{\max}/2$ and $\tilde{s}_{\min}/2 > c$. Note that first-best industry profits are given by

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

PROPOSITION 2. Let δ^* (δ^{**}) be the cutoff discount factor above which first-best industry profits can be sustained in equilibrium when firms observe only industry profits can be sustained in equilibrium when firms observe only industry profits demand (observe all prices and sales). If $\phi \neq 1/2$ (so that market demand exhibits either positive or negative persistence), then $\delta^* < \delta^{**}$. In particular, if $\delta \in (\delta^*, \delta^{**})$, then first-best industry profits can be sustained if firms observe only industry sales, but not if firms also observe each of their competitor's prices and sales.

The intuition is simple. Under the home-market principle, when $\phi > 1/2$, 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 $\phi < 1/2$ case is symmetric). The home-market principle is therefore sustainable if and only if this deviation is deterred by the threat of reverting to 0 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 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 just below

¹⁹ The extreme assumption that there is no aggregate demand uncertainty whatsoever could easily be relaxed. As becomes clear below, all that is really needed is that a firm cannot invert aggregate demand to determine the full vector of market-level demand states. Note that this property does imply some degree of negative correlation in the support of the vector of market-level demand states, because if demand states are independent or positively correlated, then, when aggregate demand takes on its highest possible level, a firm can infer that demand in each market must also be at its maximum.

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

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

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

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

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.

Figure 1 illustrates this intuition. When $\phi \neq 1/2$, the cutoff discount factor is strictly higher when all prices and quantities are observable (“perfect monitoring”). The cutoff δ^{**} is U-shaped in this case, since a more extreme value of ϕ generates more extreme beliefs about on-path continuation payoffs and thus more extreme pessimism about on-path continuation payoffs at those histories where firms are most tempted to deviate.

In contrast, when firms observe only industry demand (“imperfect monitoring”), the cutoff discount factor δ^* has two U-shaped regions, between




FIG. 1.—Comparison of cutoff discount factors δ^{**} (perfect monitoring) and δ^* (imperfect monitoring) when $c = 1$, $s_H = 4$, $s_L = 3$, and $n = 10$. Color version available as an online enhancement.

$[\phi = 0, \phi = 1/2]$ and $[\phi = 1/2, \phi = 1]$. Focusing, for example, on the region $[\phi = 1/2, \phi = 1]$, the intuition is that, when ϕ is close to $1/2$, the best deviation price is $\tilde{s}_L/2$, so the best deviation profit is decreasing in ϕ ; at some point the best deviation price switches to $\tilde{s}_H/2$, after which the best deviation profit is increasing in ϕ .

We have 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 firm. (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 each firm can observe its competitors' prices, in addition to its 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 market demand states in period $t - 2$. If a 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

$$\frac{[\tilde{s}_H - (2\phi - 1)(\tilde{s}_H - \tilde{s}_L)](\tilde{s}_L - 2c)}{8} + \phi \frac{\tilde{s}_H(\tilde{s}_H - 2c)}{8}.$$

A firm's maximum deviation gain in this model (when $\phi > 1/2$) is therefore

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

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.

B. Example 2: Better Detection versus Better Deviations

Consider the same setting as in example 1, but with only two firms and markets and with demand state transitions independent across markets. Each demand state continues to have persistence ϕ , which we assume is greater than $1/2$. We are again interested in the minimum discount factor required to sustain the outcome where each firm prices at $\tilde{s}_L/2$ or $\tilde{s}_H/2$ when the previous home-market state is low or high, respectively. Assume that $\tilde{s}_H(\tilde{s}_L - 2c) \leq \tilde{s}_H(\tilde{s}_H - 2c)/2$, so that in the absence of information about the foreign-market demand state the most tempting deviation price is (just below) $\tilde{s}_H/2$.

Suppose that the precision of the monitoring technology is indexed by $\rho \in [0, 1]$ and that ρ enters the model in two ways:

1. Firms receive binary signals of the foreign-market demand state, and the precision of this signal is $(1 + \rho)/2$.²⁰
2. If a firm enters its competitor's market, a public signal revealing this fact realizes with probability ρ .

Formally, we may write $z_i = (z_i^1, z_i^2, z_i^3) \in \{L, H\} \times \{In_i, Out_i\} \times \{In_j, Out_j\}$, where the distributions over the three components of z_i are conditionally independent and satisfy $\Pr(z_i^1 = L | s^j = s_L) = \Pr(z_i^1 = H | s^j = s_H) = (1 + \rho)/2$, $\Pr(z_i^2 = In_i | p_i^j = \infty) = \Pr(z_i^3 = In_j | p_j^i = \infty) = 0$, $\Pr(z_i^2 = In_i | p_i^j < \infty) = \Pr(z_i^3 = In_j | p_j^i < \infty) = \rho$, and $\Pr(z_i^2 = z_j^2) = \Pr(z_i^3 = z_j^3) = 1$.²¹ Note that only the case $\rho = 1$ satisfies assumption 1 (as this entails perfect detection of entry into one's home market), while $\rho = 0$ corresponds to the case where firms obtain no information beyond their own sales and intermediate values of ρ interpolate linearly between these two extremes. In particular, both effect M and effect D are present in this example.

We claim that the cutoff discount factor required for sustaining the home-market principle is the same at $\rho = 0$ and $\rho = 1$, while it is strictly lower at any ρ between 0 and 1. The interplay of effects M and D thus makes an intermediate level of transparency optimal.

²⁰ Under the home-market principle where the foreign firm only ever prices at $\tilde{s}_L/2$ or $\tilde{s}_H/2$, this is equivalent to assuming that firms receive binary signals of their competitor's sales, namely, a signal of whether one's competitor's sales lies in the set $\{s^H - \tilde{s}_L/2, s^H - \tilde{s}_H/2\}$ or the set $\{s^L - \tilde{s}_L/2, s^L - \tilde{s}_H/2\}$.

²¹ We maintain the assumption that each firm observes its own past prices and sales and its own current costs. Thus, consistent with the notation throughout the paper, the signal z represents a firm's additional information beyond these variables.

To see this, note that the key incentive constraint under which the home-market principle is an equilibrium is

$$\delta V_L \geq \psi \frac{\tilde{s}_H(\tilde{s}_H - 2c)}{4} + \delta(1 - \psi)(1 - \rho)V_L,$$

where V_L is a firm's continuation payoff under the home-market principle when its previous home-market demand state was low (note that this is independent of ρ) and ψ is the greatest probability that firm i can ever assign (along the equilibrium path) at the beginning of period t to the event that firm j prices at $\tilde{s}_H/2$ in period t . This follows because entering the foreign market with price $\tilde{s}_H/2$ both yields a greater instantaneous payoff than does entering with any other price and minimizes the probability of detection among all deviations with a positive probability of making sales.

Rewriting the incentive constraint as

$$V_L \geq \frac{1}{\delta} \frac{1}{1 + \rho[(1/\psi(\rho)) - 1]} \frac{\tilde{s}_H(\tilde{s}_H - 2c)}{4},$$

where we have made explicit the dependence of ψ on ρ , we obtain the following result.

PROPOSITION 3. The cutoff discount factor above which first-best industry profits can be sustained in equilibrium depends on ρ only through the quantity $\rho[(1/\psi(\rho)) - 1]$ and is decreasing in this quantity. Thus, the greater is $\rho[(1/\psi(\rho)) - 1]$, the easier it is to sustain collusion.

Furthermore, $\rho[(1/\psi(\rho)) - 1]$ is inverse-U shaped in ρ and equals 0 if $\rho = 0$ or $\rho = 1$, and the value of ρ that maximizes $\rho[(1/\psi(\rho)) - 1]$ is decreasing in ϕ . Thus, collusion is easier to sustain when the precision of monitoring is intermediate, and the level of precision of monitoring that makes sustaining collusion easiest is decreasing in the persistence of demand.

The intuition for proposition 3 is as follows.

If $\rho = 1$, then each firm perfectly observes the previous foreign demand state. The most tempting deviation is thus to wait until the previous home demand state is low and the previous foreign demand state is high, and then enter the foreign market with price $\tilde{s}_H/2$. This deviation is unprofitable if and only if the equilibrium continuation payoff δV_L exceeds the instantaneous gain from stealing the foreign market $\tilde{s}_H(\tilde{s}_H - 2c)/4$.

If $\rho = 0$, firms have no information about the foreign demand state but also are not detected when they enter the foreign market unless they successfully steal the market. A firm can therefore wait until its home demand state is low and then enter the foreign market with price $\tilde{s}_H/2$: this deviation is detected if and only if it successfully steals the foreign market, so the relevant incentive constraint is again that δV_L must exceed $\tilde{s}_H(\tilde{s}_H - 2c)/4$, exactly as in the $\rho = 1$ case.

However, if ρ is strictly between 0 and 1, then the most tempting deviation is not as attractive as in the $\rho = 1$ case or the $\rho = 0$ case. In particular, the best deviation is now to wait until the previous home demand state is low and the belief about the previous foreign demand state is as optimistic as possible: that is, until $\Pr(s_{t-1}^j = H)$ is very close to $\psi(\rho)$. Since $\psi(\rho)$ is strictly less than 1 for all $\rho < 1$ (as verified in the proof of proposition 3) and the detection probability is strictly positive, this deviation runs a risk of being detected while failing to steal the foreign market. In contrast, this risk is entirely absent when $\rho \in \{0, 1\}$, which explains why sustaining collusion is easier when ρ is intermediate.

In terms of the effects discussed in Section I, the issue is that potential deviations are very profitable when ρ is high (effect D) and are very hard to detect when ρ is low (effect M). Collusion is thus easiest to sustain when ρ is intermediate.

Finally, the intuition for the result that optimal monitoring precision is decreasing in demand persistence may be seen by considering the extreme cases. If demand is independently and identically distributed (i.i.d.), then effect D is absent, so $\rho = 1$ is optimal, while if demand is very persistent, then uncertainty about demand can be preserved only if ρ is close to 0.

C. Example 3: Better On-Path Pricing versus Better Deviations

The next example retains assumption 1 while relaxing assumption 3 and thus examines the trade-off between effects C and D.

There are two firms, two markets, and two market demand states, s_L and s_H . Period t demand in the two markets is perfectly correlated and is also perfectly correlated with a random variable s_{t-1}^3 that realizes in period $t - 1$. (The notation indicates that this random variable can be interpreted as the period $t - 1$ demand state in a third, “dummy” market, in which it is infeasible for the firms to sell because of trade or technological barriers.) The state s_t^3 is i.i.d. across periods with $\Pr(s_t^3 = s_L) = \Pr(s_t^3 = s_H) = 1/2$. Demand in each market $j = 1, 2$ is again linear: $D(p^j, s^j) = s^j - p^j$. Each firm can produce at zero cost in either market.²² Assume that $\delta \geq 1/2$.

At the beginning of period t , the firms observe a common signal z_t of the state s_{t-1}^3 , with $z_t \in \{z_L, z_H\}$ and

$$\Pr(z = z_L | s = s_L) = \Pr(z = z_H | s = s_H) = \frac{1 + \rho}{2},$$

²² 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.

where $\rho \in [0, 1]$ again measures the precision of the signal. For instance, if $\rho = 1$, then at the end of period t the firms perfectly learn the period $t + 1$ demand state in both markets, while if $\rho = 1/2$, they obtain no information about the period $t + 1$ demand state. In this example, let $\tilde{s}_L = \rho s_L + (1 - \rho)s_H$ and let $\tilde{s}_H = (1 - \rho)s_L + \rho s_H$.

This information structure violates assumption 3, as (for instance) s_t^1 is not independent of s_{t-1}^3 , conditional on s_{t-1}^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 precision ρ that allows for the greatest industry profits and 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.²³ As s_t —and therefore z_t —are i.i.d. across periods with equal probability on each realization, a firm's sequential rationality conditions after signals z_L and z_H are, respectively,

$$(1 - \delta)u_L \leq \delta \left(\frac{u_L}{2} + \frac{u_H}{2} \right), \quad \text{and}$$

$$(1 - \delta)u_H \leq \delta \left(\frac{u_L}{2} + \frac{u_H}{2} \right).$$

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

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

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-period, per-market) industry profits equal

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

²³ 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.

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

$$\delta^*(1) = \frac{2}{3 + (\sigma_L/\sigma_H)^2}.$$

There are two cases:

1. If $\delta < \delta^*(\rho)$, then $V^* = (1/2)\{1 + [\delta/(2 - 3\delta)]\}(\tilde{\sigma}_L^2/4)$, and V^* is increasing in δ and decreasing in ρ .
2. If $\delta > \delta^*(\rho)$, then $V^* = (1/2)[(\tilde{\sigma}_L^2/4) + (\tilde{\sigma}_H^2/4)]$, and V^* is constant in δ and increasing in ρ .

We can now read off the optimal value of ρ . Let $\rho^*(\delta) = \{\rho : \delta^*(\rho) = \delta\}$, and note that $\rho^*(\delta)$ is an increasing function on $\delta \in [1/2, \delta^*(1)]$.²⁴

PROPOSITION 4. The level of precision ρ that maximizes industry profits V^* is given by $\rho = \min\{\rho^*(\delta), 1\}$. Thus, the optimal level of precision is increasing in δ , and it lies strictly between 0 and 1 when $\delta \in (1/2, \delta^*(1))$.

Proof. The proposition follows because V^* is increasing in ρ if $\rho < \rho^*(\delta)$ and decreasing in ρ if $\rho > \rho^*(\delta)$. QED

To understand this result, note that giving the firms more information about the state always increases unconditional expected first-best profits, $(\tilde{\sigma}_L^2 + \tilde{\sigma}_H^2)/8$ but decreases expected first-best profits after the bad signal, $\tilde{\sigma}_L^2/4$. 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 $\tilde{\sigma}_L^2/4$.²⁵ Therefore, if δ is close to 1/2, then the optimal information structure is less informative (to maximize $\tilde{\sigma}_L^2/4$), while if δ is close to 1, then the optimal information structure is more informative (to maximize $(\tilde{\sigma}_L^2 + \tilde{\sigma}_H^2)/8$).²⁶

A more general intuition is simply that providing more precise information about demand increases collusive profits by more than it increases the gain from deviating if and only if firms are sufficiently patient.

²⁴ 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.

²⁵ This is also an implication of proposition 4 of Kandori (1991b).

²⁶ 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.

D. *Example 4: Transparency Can Hinder Collusion
in a Stationary Environment*

In the analysis so far, optimal collusive equilibria were stochastic for the obvious reason that the underlying physical environment was assumed to be stochastic. As we have seen, this stochasticity opens up the possibility that a transparent information structure can help firms devise more profitable deviations and can thereby hinder collusion. In our last example, we show that—perhaps surprisingly—transparency can also hinder collusion even when the physical environment is completely stationary and free from uncertainty. The intuition is that the need to provide intertemporal incentives alone can lead the cartel to follow a stochastic equilibrium, which again creates the possibility that transparency can hinder collusion.

Specifically, consider the special case of the model in Section III where demand states and production costs are known and fixed over time. Assume that there are two firms and two markets, with $c_1^1 = c_2^2 = 0$, $c_1^2 = \infty$, and $c_2^1 = 0.7$, and assume unit demand in each market with a choke price of 1: $D(p) = 1$ if $p \leq 1$, and $D(p) = 0$ if $p > 1$. (Note that we suppress the demand state s , as it is now held constant.) Thus, selling in market 2 is prohibitively costly for firm 1, but selling in market 1 is potentially profitable for firm 2.

We refer to any outcome of the game where in every period each firm sets a finite price only in its home market as *market segmentation*. Note that productive efficiency requires that each firm makes sales only in its home market, while market segmentation strengthens this by also requiring that firms do not enter the foreign market even at prices that do not result in sales. Market segmentation would thus be necessary for efficiency if there were a small fixed cost of entering the foreign market.²⁷

We will show that there is a range of discount factors Δ such that, for any $\delta \in \Delta$, the following results hold.

1. If prices are perfectly observable, then market segmentation is not sustainable, even if the firms can rely on an intermediary to help them correlate their actions.
2. If each firm observes only its own price and sales, then market segmentation is sustainable with the assistance of an intermediary.

To state these results formally, we need to be more specific about the role of the intermediary, henceforth referred to as a *mediator*. We assume that the mediator can perfectly observe the firms' past actions and can communicate privately with the firms. A standard application of the revelation principle (Forges 1986; Myerson 1986; Sugaya and Wolitzky 2018b)

²⁷ All incentive constraints in the analysis that follows are strict, so our results are robust to explicitly introducing such a fixed cost into the model.

implies that in such games there is no loss of generality in restricting attention to so-called *obedient equilibria*, where in each period t the game proceeds as follows.²⁸

1. The mediator makes a private action recommendation $r_{i,t} \in A_i$ to each firm $i = 1, 2$, where A_i is firm i 's set of available actions (i.e., prices in each market). This recommendation may be conditioned on the mediator's private history, $h^t = (r_\tau, a_\tau)_{\tau=0}^{t-1}$, which consists of the entire history of both recommendations and actions.
2. Each firm i takes an action $a_{i,t} \in A_i$. This action may be conditioned on firm i 's private history, $h_i^t = (r_{i,\tau}, z_{i,\tau})_{\tau=0}^{t-1}$, which consists of the history of private recommendations to firm i and firm i 's private signals. By definition of an obedient equilibrium, firm i 's equilibrium strategy is to obey the mediator's recommendation (i.e., play $a_{i,t} = r_{i,t}$).
3. The mediator observes the realized action profile a_t . However, the distribution of a firm's signal $z_{i,t}$ depends on the information structure. With perfect monitoring, $z_{i,t} = a_t$, so actions (i.e., prices) are observable. With private monitoring, $z_{i,t} = (a_{i,t}, u_{i,t})$, so each firm observes only its own action and its realized payoff.

We establish the following result in appendix C.

PROPOSITION 5. There is an open interval of discount factors Δ such that, for any $\delta \in \Delta$, market segmentation is not sustainable with mediated perfect monitoring but is sustainable with mediated private monitoring.

Before discussing the reasoning behind this result, let us comment on the realism of the two roles played by the mediator in the model: communicating correlated information to the firms and monitoring the firms' actions.

The presence of a mediator who communicates correlated information to the firms seems quite realistic in light of the examples discussed in Sections I and V: many real-world cartels do rely on intermediaries to help them collude, and the intermediaries' roles often involve summarizing (and thus coarsening) more detailed information about cartel participants' behavior.²⁹ In addition, 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,

²⁸ For a more detailed exposition of repeated games with a mediator, see Sugaya and Wolitzky (2017).

²⁹ Levenstein and Suslow (2006, 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."

a “message” of the mediator’s would instead be interpreted as a firm’s signal of its competitors’ behavior. In Sugaya and Wolitzky (2017), we discuss what properties such a private monitoring structure would need to have to justify this interpretation.³⁰

The assumption that the mediator directly observes the firms’ actions—and in particular does not need to rely on self-reports by the firms—may also 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, in an earlier version of this paper we have also shown that the assumption that the mediator directly observes actions can be completely dispensed with if firm profits are noisy, albeit at the cost of a more complicated equilibrium construction.

Now, what is the intuition for why transparency can prevent market segmentation in the current example? This is more easily explained in a simplified setting where each firm has only two price levels available: a low price that yields lower profits but serves to deter entry into one’s home market and a high price that yields higher profits but encourages entry.³¹ In this setting, for sufficiently low discount factors, firm 1 is so impatient that she is willing to price low only if she is rewarded by pricing high while firm 2 stays out of her market in the very next period. With perfect monitoring, this means that, when firm 1 prices low in period t , firm 2 observes this action and then knows to expect a high price in period $t + 1$. But firm 2 is too impatient to stay out of firm 1’s market when firm 1 prices high, so he will enter. This implies that firm 1 will never price low. Hence, market segmentation cannot be maintained under perfect monitoring.

With private monitoring, however, the mediator can recommend that firm 1 alternate between high and low prices without informing firm 2 of whether firm 1 prices high in even or odd periods. Firm 2 therefore never expects firm 1 to price high with probability greater than 50 percent and is thus willing to stay out of firm 1’s market. Meanwhile, as firm 1 always prices high in period $t + 1$ after pricing low in period t , she receives the intertemporal reward required to make pricing low incentive compatible. This arrangement therefore succeeds in segmenting the market.

Concisely put, in this example market segmentation requires a stochastic intertemporal dependence between high and low prices, and along such a stochastic path of play revealing a firm’s past actions will prompt a deviation by its competitor.

³⁰ 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.

³¹ This simplified game is thus a multimarket version of a standard entry-deterrence game, as in Selten (1978).

Let us comment on the range of discount factors Δ for which proposition 5 holds. The range of discount factors allowed in the proof is (0.148, 0.149). One would ideally want a result that holds for a wider range of discount factors and for higher discount factors. The range of discount factors could be widened by constructing a more complicated equilibrium under private monitoring; our proof is optimized for simplicity, not for the size of Δ . We can also allow one of the firms to have a much higher discount factor if we consider asymmetric market demand and heterogeneous discounting; specifically, we show in appendix C that if demand in market 1 is scaled up by a factor of 100, then the conclusion of proposition 5 holds for $\delta_1 = 0.142$ and $\delta_2 = 0.949$.

V. 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 our results. 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.³² The goal is only to connect our theory with some observed cartels and to highlight the theory's plausibility.

A. Antitrust Cases

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

1. 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

³² For example, an obvious alternative explanation for coarsening information is that this may make firms more willing to reveal their private information. While this is certainly plausible, an advantage of our theory is that it remains valid even if the cartel is able to directly audit its members' books, which, as we have noted, is often a possibility.

other end consumers” (Harrington 2006, 85; also see EC 2004, 10–11). From 1988 to 2001, the European copper plumbing tubes industry—a roughly €1 billion industry—was cartelized by a group of five to nine firms that jointly accounted for approximately 65–80 percent of the market (EC 2004, 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—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’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” (EC 2004, 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” (EC 2004, 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” (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” (75). It thus appears that the cartel shifted from sharing firm-level data to sharing 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” (EC 2004, 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” (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” (115).

Finally, there is also evidence that detecting entry into one’s home market was particularly easy in the copper plumbing tube industry. According to the EC, “At least until 1995, monitoring was facilitated by national certification procedures. Copper plumbing tubes had to be certified in each Member State. Each Member State had its own certification label. Certification organisations . . . prohibited producers at least until 1995 to indicate different national certifications on plumbing tubes” (EC 2004, 35).

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.

2. Isostatic Graphite

Isostatic graphite is a graphite product used in industrial applications such as the production of certain types of electrodes and semiconductors (EC 2002a, 6). The EC prosecuted eight firms for cartelizing the European isostatic graphite industry (a roughly €500 million industry) in the mid-1990s. 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” (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” (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” (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.

3. Methylglucamine

Methylglucamine is “an intermediate chemical product for the synthesis of x-ray media, pharmaceuticals and colourings” (EC 2002c, 21). Throughout the 1990s, the only two producers of pharmaceutical-grade methyl-

glucamine in the world were Merck and Aventis/Rhône-Poulenc Biochimie (RPB; 23). The EC found that these two firms “formed a clandestine cartel . . . by which they fixed market shares . . . ; agreed on price targets . . . ; agreed on price lists . . . ; and agreed on how to share the largest customers” (20). Most relevant for our model is that both the home-market principle and imperfect information about market shares seem to have played important roles in the methylglucamine cartel. In particular, “the parties agreed not to compete for the other party’s customers” (30), and therefore “the Commission consider[ed] that it is established that the parties agreed to share the market through customer allocation” (29). While “oral exchange of sales figures did occur,” this “did not materialise into a full systematic exchange of sales data” (28), with the consequence that “both producers . . . inaccurately assessed one another’s position in the meglumine [i.e., methylglucamine] market. Merck mistakenly believed that both producers had a 50% market share. . . . [But] Merck had a much higher market share than RPB, i.e., Merck had [around 65%] of the world market. . . . In turn, RPB . . . underestimated the worldwide meglumine market” (28). This pattern of (1) market segmentation through customer allocation combined with (2) imperfect information about market shares, resulting from limited information exchange, is again consistent with our model.

B. 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 chloride, district heating pipes, electrical and mechanical carbon graphite, lysine, methionine, nucleotides, seamless steel tubes, soda ash, vitamins, and zinc phosphate (Harrington 2006, 34–40), in addition to the copper plumbing tubes and isostatic graphite cartels discussed above. 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 to have played an important role in the cartels in plasterboard (Harrington 2006, 54) and low-density polyethylene (Marshall and Marx 2012, 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” (EC 2002b, 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” (6).

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

C. Suggestive Correlational Evidence

Finally—while we do not wish to overemphasize this point, given the sparseness of the available data—there may even be a slight statistical correlation between cartels’ use of information coarsening and their reliance on the home-market principle. Of the four cartels that seemed to rely heavily on information coarsening (copper plumbing tubes, isostatic graphite, plasterboard, and low-density polyethylene), all but one (low-density polyethylene) appear to have also relied on the home-market principle; while of the 20 other cartels discussed in Harrington’s survey, only the 10 cartels referenced above (“choline chloride . . . zinc phosphate”) are cited as following the home-market principle.

There also may be a correlation between cartels’ reliance on intermediaries to manage their informational environment (whether to coarsen information or not) and reliance on the home-market principle. In table 6.1 of their book, Marshall and Marx (2012, 126–27) classify the 22 major industrial cartel decisions of the EC from 2000 through 2005 according to whether the cartel relied on customer, geographic, and/or market share allocation and whether the cartel relied on a third-party facilitator. According to their classification, 19 of the 22 cartels used some form of allocation scheme, and 12 used customer or geographic allocation (corresponding to the home-market principle). Of the 19 cartels that used some allocation scheme, 11 also relied on a third-party facilitator; while of the 12 that used customer or geographic allocation, seven relied on a third party. In contrast, none of the three cartels that did not use an allocation scheme relied on a third party.

In addition, in an insightful discussion of this paper, Leslie Marx has observed that various subsets of a group of nine of the European chemicals firms were involved in seven distinct cartels, three of which were facilitated by the Swiss consulting firm AC-Treuhand. She reports that two of the three Treuhand-facilitated cartels relied on both geographic allocation and customer allocation, while three of the four non-Treuhand cartels relied on geographic allocation but only one relied on customer allocation.

Thus, while the available correlational evidence is weak (e.g., none of the above relationships are statistically significant) and even the under-

lying classifications are highly subjective, the evidence at least seems to point in the direction suggested by the theory.

VI. Conclusion

The goal of this paper has been a reassessment of Stigler's pathbreaking idea that transparency within a cartel facilitates collusion. In contrast to this idea, we find that—under some assumptions—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—and have derived testable comparative statics predictions—by considering parameterized examples featuring both costs and benefits of transparency. And we have further shown that transparency can hinder collusion even in a stationary economic environment.

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 that is quite different from existing approaches (Athey, Bagwell, and Sanchirico 2004; Harrington and Chen 2006). Consider the following example. 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 one's current home demand state—and hence revealing information about one's future home 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 colluding firms can obtain by keeping their actions and outcomes 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 Theorem 1

In what follows, all omitted proofs of lemmas may be found in appendix C.

Let $\Delta(\pi)$ denote the set of discount factors for which the home-market principle is sustainable in perfect Bayesian equilibrium with information structure π .

Let $I_i^t := (c_t, s_{t-1}^i, y_{i,t}, (x_{j,i,t}^j)_{j \neq i})_{\tau=0}^t$ denote the vector of costs; past own-market demand states; and signals of sales, costs, and other firms' home-market prices for firm i up to the beginning of period t . Under assumptions 2 and 4, this information is available to firm i at the beginning of period t under any strategy profile satisfying the home-market principle. Let $\mathcal{I}_i^t(\pi)$ denote the set of vectors I_i^t that arise with positive probability under the home-market principle with information structure π , and let $\mathcal{I}_i(\pi) = \cup_t \mathcal{I}_i^t(\pi)$; note that these sets are the same for all strategy profiles satisfying the home-market principle (i.e., they do not depend on which losing prices firms set in foreign markets).

As a firm's minmax payoff is 0, a necessary condition for the home-market principle to be sustainable in equilibrium is that, for each firm i , each period t , and each $I_i^t \in \mathcal{I}_i^t(\pi)$,

$$\begin{aligned} & \mathbb{E} \left[\sum_{\tau \geq t+1} \delta^{\tau-t} D(p_i^m(c_{i,\tau}^i, s_{\tau-1}^i), s_\tau^i)(p_i^m(c_{i,\tau}^i, s_{\tau-1}^i) - c_{i,\tau}^i) | I_i^t \right] \\ & \geq \max_{(p_{i,t}^j)_{j \neq i}} \sum_{j \neq i} \mathbb{E} \left[\mathbf{1}_{\{p_{i,t}^j \geq p_{i,t}^i\}} D(p_{i,t}^j, s_{i,t}^j)(p_{i,t}^j - c_{i,t}^j) | I_i^t \right], \end{aligned} \tag{A1}$$

where $\mathbf{1}_{\{ \cdot \}}$ denotes the indicator function.

The left-hand side of equation (A1) depends on I_i^t only through the pair $(c_{i,t}^i, s_{i,t-1}^i)$.

LEMMA 1. For each i and δ , there exists a function $v_i^\delta : C_i^i \times S^i \rightarrow \mathbb{R}$ such that, for each $I_i^t \in \mathcal{I}_i(\pi)$,

$$\mathbb{E} \left[\sum_{\tau \geq t+1} \delta^{\tau-t} D(p_i^m(c_{i,\tau}^i, s_{\tau-1}^i), s_\tau^i)(p_i^m(c_{i,\tau}^i, s_{\tau-1}^i) - c_{i,\tau}^i) | I_i^t \right] = v_i^\delta(c_{i,t}^i, s_{i,t-1}^i).$$

Furthermore, $v_i^\delta(c_i^j, s^j)$ is increasing in δ for all $(c_i^j, s^j) \in C_i^j \times S^j$.

So far, equation (A1) is only a necessary condition, as it considers a firm's incentives conditional on only the information contained in I_i^t , rather than on the firm's full information set (which also contains the signals of other firms' foreign-market prices, $(x_{j,i,t}^k)_{\tau=0}^t$ with $j \neq k$). It may be shown to be sufficient by constructing standard "grim trigger" strategies, where any firm that enters a foreign market

or detects entry into its home market prices at $p_{i,t}^j = c_{j,t}^i$ in every market j in every subsequent period t .

LEMMA 2. $\delta \in \Delta(\pi)$ if and only if, for each i and each $I_i^t \in \mathcal{I}_i(\pi)$,

$$v_i^\delta(c_{i,t}^i, s_{i-1}^i) \geq \max_{(p_{i,t}^j)_{j \neq i}} \sum_{j \neq i} \mathbb{E} \left[\mathbf{1}_{\{p_{i,t}^j(c_{j,t}^i, s_{i-1}^i) \geq p_{i,t}^j\}} D(p_{i,t}^j, s_{i,t}^j) (p_{i,t}^j - c_{i,t}^i) | I_i^t \right]. \tag{A2}$$

Combining lemmas 1 and 2, we see that $\Delta(\pi) = [\delta^*(\pi), 1]$, where $\delta^*(\pi) \in (0, 1)$ is the unique solution to

$$\sup_{I_i^t \in \mathcal{I}_i(\pi)} \max_{(p_{i,t}^j)_{j \neq i}} \sum_{j \neq i} \mathbb{E} \left[\mathbf{1}_{\{p_{i,t}^j(c_{j,t}^i, s_{i-1}^i) \geq p_{i,t}^j\}} D(p_{i,t}^j, s_{i,t}^j) (p_{i,t}^j - c_{i,t}^i) | I_i^t \right] - v_i^\delta(c_{i,t}^i, s_{i-1}^i) = 0. \tag{A3}$$

The remainder of the proof thus consists of showing that the first term of equation (A3) (the ‘‘maximum deviation gain’’) is always at least weakly increasing in π in the Blackwell order and is strictly increasing in π in the Blackwell order under assumptions 5 and 6.

We begin by rewriting the maximum deviation gain as a function of the set of beliefs over demand states that may arise in equilibrium. Formally, let $b_i(I_i^t)$ be the distribution over states S^{-i} conditional on I_i^t under the home-market principle, and let $B_i(\pi)|_{c,s^i} = \{b_i(I_i^t) : I_i^t|_{(c,s_{i-1}^i)} = (c, s^i), I_i^t \in \mathcal{I}_i(\pi)\}$. In addition, let

$$\Pi(p_i^{-i}, c, s^{-i}) := \sum_{j \neq i} \mathbb{E} \left[\mathbf{1}_{\{p_{i,t}^j(c_{j,t}^i, s_{i-1}^i) \geq p_{i,t}^j\}} D(p_{i,t}^j, s_{i,t}^j) (p_{i,t}^j - c_{i,t}^i) | c, s^i \right]$$

be firm i 's profit from setting prices p_i^{-i} in the foreign markets at cost vector c when the previous foreign market demand state is s^{-i} , and let $\Pi(p_i^{-i}, c, b_i) := \sum_{s^{-i}} b_i(s^{-i}) \Pi(p_i^{-i}, c, s^{-i})$ be firm i 's expected profit from setting prices p_i^{-i} in the foreign markets at cost vector c and belief b_i . Finally, denote firm i 's maximum deviation gain at cost vector c and belief b_i by $d_i(c, b_i) := \max_{p_i^{-i}} \Pi(p_i^{-i}, c, b_i)$. Note that $\Pi(p_i^{-i}, c, b_i)$ is linear in b_i , so d_i is the upper envelope of linear functions of b_i and is therefore convex in b_i . Note also that the first term of equation (A3) equals $\max_{(c,s^i) \in C \times S} \sup_{b_i \in B_i(\pi)|_{c,s^i}} d_i(c, b_i)$. Thus, it remains to show that $\max_{(c,s^i) \in C \times S} \sup_{b_i \in B_i(\pi)|_{c,s^i}} d_i(c, b_i)$ is increasing in π in the Blackwell order.

To see this, consider the following degenerate auxiliary game, with $n + 1$ players and no actions. Player 0 (who corresponds to nature in the original game) has initial state $(c_0, s_0) \in C \times S$ drawn according to \mathfrak{q} , and her state transitions according to M . At the beginning of period t , each player $i \neq 0$ observes the signal $(s_{i-1}^i, c_i, z_{i,t})$,³³ where z_i is distributed according to $\pi^m(z_i | s_{i-2}, c_{i-1}, s_{i-1}, c_i) := \pi(z_i | (p_{i,t}^m(c_{i,t-1}^i, s_{i-2}^i)_{i=1}^n, s_{i-1}, c_i))$. We say that player i 's period t state is (c_i, s_{i-1}^i) , and let $b_i|_{c_i, s_{i-1}^i}(z_{i,t})$ be player i 's belief about s_{i-1} when her state is (c_i, s_{i-1}^i) and she receives signal $z_{i,t}$.

We wish to characterize the set of beliefs $b_i \in \Delta(S^{-i})$ that can arise in equilibrium in this game.³⁴ Following Phelan and Skrzypacz (2012; cited as PS in the online appendix), for each i , define the mapping $T_{i,\pi}^U : \Delta(S^{-i})^{C \times S^i} \rightarrow \Delta(S^{-i})^{C \times S^i}$ by

$$T_{i,\pi}^U((B_i|_{c,s^i})_{c,s^i}) = \text{co}((B_i|_{c,s^i})_{c,s^i} \cup T_{i,\pi}((B_i|_{c,s^i})_{c,s^i})),$$

³³ Except in period 0, where each player observes only c_0 .

³⁴ In what follows, $\text{co}(\cdot)$, $\text{cl}(\cdot)$, $\text{int}(\cdot)$, and $\text{ext}(\cdot)$ stand for convex hull, closure, interior, relative interior, and extreme points, respectively.

where

$$T_{i,\pi}((B_i|_{c,s^i})_{c,s^i}) = \left\{ \left\{ \begin{array}{l} (b^i|_{c,s^i})_{c,s^i} \in (\Delta(S^{-i}))_{c,s^i} : \\ \exists (\hat{c}, \hat{s}^i), (b|_{\hat{c},\hat{s}^i}) \in B_i|_{\hat{c},\hat{s}^i}, z_i \\ \text{such that } b^i|_{c,s^i} = B_{i,\pi}^{c,s^i}(b|_{\hat{c},\hat{s}^i}, z_i) \text{ for each } c, s^i \end{array} \right\} \right\},$$

and³⁵

$$\mathcal{B}_{i,\pi}^{c,s^i}(b|_{\hat{c},\hat{s}^i}, z_i) = \left(\frac{\sum_{\hat{s}^{-i}} b_{\hat{c},\hat{s}^i}(\hat{s}^{-i})M(c, s|\hat{c}, \hat{s})\pi^m(z_i|\hat{s}, \hat{c}, s, c)}{\sum_{\hat{s}^{-i}} \sum_{\hat{s}^{-i}} b_{\hat{c},\hat{s}^i}(\hat{s}^{-i})M(c, s^i, \hat{s}^{-i}|\hat{c}, \hat{s})\pi^m(z_i|\hat{s}, \hat{c}, s^i, \hat{s}^{-i}, c)} \right)_{\hat{s}^{-i}}.$$

Thus, $T_{i,\pi}((B_i|_{c,s^i})_{c,s^i})$ is the set of possible posterior beliefs at each state (c, s^i) when the set of possible prior beliefs at each state (c, s^i) is $B_i|_{c,s^i}$. Finally, $T_{i,\pi}^U((B_i|_{c,s^i})_{c,s^i})$ is the set of possible beliefs that can arise as convex combinations of prior and posterior beliefs.

Lemma 2 and section 3.1 of Phelan and Skrzypacz (2012) establish the following facts. For all i and c_0 , let $B_i|_{c_0,s_{-1}} = \{\varrho\}$.³⁶ Then:

1. There exists a smallest fixed point of the operator $T_{i,\pi}^U$ containing ϱ . Denote this fixed point $\mathcal{M}(T_{i,\pi}^U)$.
2. $\mathcal{M}(T_{i,\pi}^U)|_{c,s^i}$ is the closure of the set of on-path beliefs consistent with state (c, s^i) : $\mathcal{M}(T_{i,\pi}^U)|_{c,s^i} = \text{cl}(B_i(\pi)|_{c,s^i})$.
3. $\mathcal{M}(T_{i,\pi}^U) = \text{cl}(\lim_{k \rightarrow \infty} (T_{i,\pi}^U)^k(B_i|_{c_0,s_{-1}}))$.
4. $\mathcal{M}(T_{i,\pi}^U)$ is a compact and convex subset of $\Delta(S^{-i})^{C \times S^i}$.
5. If $\varrho \in \text{int}(\mathcal{M}(T_{i,\pi}^U))$, then for every extreme point b_i of $\mathcal{M}(T_{i,\pi}^U)$ and every (c, s^i) , there exists $\hat{b}_i \in \mathcal{M}(T_{i,\pi}^U)$, \hat{c} , \hat{s}^i , and z_i such that $b_i|_{c,s^i} = \mathcal{B}_{i,\pi}^{c,s^i}(\hat{b}_i|_{\hat{c},\hat{s}^i}, z_i)$.

In particular, since $\mathcal{M}(T_{i,\pi}^U)|_{c,s^i} = \text{cl}(B_i(\pi)|_{c,s^i})$, we see that $\delta^*(\pi') \geq \delta^*(\pi)$ if

$$\max_{b_i|_{c,s^i} \in \mathcal{M}(T_{i,\pi'}^U)|_{c,s^i}} d_i(c, b_i|_{c,s^i}) \geq \max_{b_i|_{c,s^i} \in \mathcal{M}(T_{i,\pi}^U)|_{c,s^i}} d_i(c, b_i|_{c,s^i}) \tag{A4}$$

for all $(c, s^i) \in C \times S^i$, and $\delta^*(\pi') > \delta^*(\pi)$ if the inequality is strict for all $(c, s^i) \in C \times S^i$.

The next lemma says that, for a single application of Bayes's rule, a more informative information structure generates more extreme beliefs and strictly so in the case of a strictly more informative information structure under a full-support assumption.

LEMMA 3. For all i , (c, s^i) , and $b|_{\hat{c},\hat{s}^i}$, if $\pi' \geq \pi$, then $\{\mathcal{B}_{i,\pi}^{c,s^i}(b|_{\hat{c},\hat{s}^i}, z_i)\}_{z_i \in Z} \subseteq \text{co}(\{\mathcal{B}_{i,\pi'}^{c,s^i}(b|_{\hat{c},\hat{s}^i}, z_i')\}_{z_i' \in Z})$. In addition, if $\pi' > \pi$; $\pi'(z|p, c, s) > 0$ for all z, p, c ,

³⁵ In words, if player i 's previous state is $(\hat{c}_{t-1}, \hat{s}_{t-2}^i)$ and her belief about s_{t-2}^{-i} is $b_{\hat{c}_{t-1},\hat{s}_{t-2}^i}$, then after receiving signal (s_{t-1}^i, c_t, z_{it}) , her belief about s_{t-1}^{-i} is $\mathcal{B}_{i,\pi}^{c_t,s_{t-1}^i}(b_{\hat{c}_{t-1},\hat{s}_{t-2}^i}, z_{it})$, by Bayes's rule.

³⁶ The notation here is simply that s_{-1} is a dummy variable introduced to maintain consistency of the notation $B_i|_{c,s_{-1}}$.

and s ; and $M(c, s|\hat{c}, \hat{s}) > 0$ for all c , s , \hat{c} , and \hat{s} , then $\{\mathcal{B}_{i,\pi}^{c,s}(b|_{i,s}, z_i)\}_{z_i \in Z} \subseteq \text{relint}(\text{co}(\{\mathcal{B}_{i,\pi}^{c,s}(b|_{i,s}, z_i)\}_{z_i \in Z}))$.

We can now complete the proof of weak monotonicity.

LEMMA 4. If $\pi' \geq \pi$, then $\delta^*(\pi') \geq \delta^*(\pi)$.

Proof. By lemma 3, $T_{i,\pi}^U(b_i) \subseteq T_{i,\pi'}^U(b_i)$ for each b_i . As $\mathcal{M}(T_{i,\pi}^U) = \text{cl}(\lim_{k \rightarrow \infty} (T_{i,\pi}^U)^k(B_i|_{c,s,i}))$, this gives $\mathcal{M}(T_{i,\pi}^U) \subseteq \mathcal{M}(T_{i,\pi'}^U)$. This implies equation (A4), and the result follows. QED

For strict monotonicity, we require another lemma. In what follows, note that assumption 6 is equivalent to $\varrho \in \text{int}(\mathcal{M}(T_{i,\pi}^U))$.

LEMMA 5. If assumptions 5 and 6 hold and $\pi' > \pi$, then for each i , $\mathcal{M}(T_{i,\pi}^U) \subseteq \text{int}(\mathcal{M}(T_{i,\pi'}^U))$.

To complete the proof, recall that $D(p^j, s^j)$ is strictly increasing in s^j . Hence, $d_i(c, b_i)$ is nonconstant on any set of beliefs $B_i \subseteq \Delta(S^{-i})$ of nonempty interior. As d_i is convex in b_i , this implies that $d_i(c, b_i)$ attains its maximum on any compact, convex set B_i only on the boundary of B_i . Hence, as $\mathcal{M}(T_{i,\pi}^U) \subseteq \text{int}(\mathcal{M}(T_{i,\pi'}^U))$, equation (A4) holds with strict inequality.

References

- Abreu, D. 1986. "Extremal Equilibria of Oligopolistic Supergames." *J. Econ. Theory* 39 (1): 191–225.
- Abreu, D., P. Milgrom, and D. Pearce. 1991. "Information and Timing in Repeated Partnerships." *Econometrica* 59 (6): 1713–33.
- Abreu, D., D. Pearce, and E. Stacchetti. 1986. "Optimal Cartel Equilibria with Imperfect Monitoring." *J. Econ. Theory* 39 (1): 251–69.
- . 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5): 1041–63.
- Albæk, S., P. Møllgaard, and P. B. Overgaard. 1997. "Government-Assisted Oligopoly Coordination? A Concrete Case." *J. Indust. Econ.* 45 (4): 429–43.
- Aoyagi, M. 2002. "Collusion in Dynamic Bertrand Oligopoly with Correlated Private Signals and Communication." *J. Econ. Theory* 102 (1): 229–48.
- . 2007. "Efficient Collusion in Repeated Auctions with Communication." *J. Econ. Theory* 134:61–92.
- Areeda, P., and L. Kaplow. 1997. *Antitrust Analysis*. 5th ed. New York: Aspen.
- Asker, J. 2010. "Leniency and Post-cartel Market Conduct: Preliminary Evidence from Parcel Tanker Shipping." *Internat. J. Indust. Org.* 28 (4): 407–14.
- Athey, S., and K. Bagwell. 2001. "Optimal Collusion with Private Information." *RAND J. Econ.* 32 (3): 428–65.
- Athey, S., K. Bagwell, and C. Sanchirico. 2004. "Collusion and Price Rigidity." *Rev. Econ. Studies* 71 (2): 317–49.
- Awaya, Y., and V. Krishna. 2016. "On Communication and Collusion." *A.E.R.* 106 (2): 285–315.
- Bagwell, K., and R. W. Staiger. 1997. "Collusion over the Business Cycle." *RAND J. Econ.* 28 (1): 82–106.
- Belleflamme, P., and F. Bloch. 2008. "Sustainable Collusion on Separate Markets." *Econ. Letters* 99 (2): 384–86.
- Bernheim, B. D., and E. Madsen. 2017. "Price Cutting and Business Stealing in Imperfect Cartels." *A.E.R.* 107 (2): 387–424.

- Bernheim, B. D., and M. D. Whinston. 1990. "Multimarket Contact and Collusive Behavior." *RAND J. Econ.* 21 (1): 1–26.
- Bhattacharjya, A., and U. B. Sinha. 2015. "Multi-market Collusion with Territorial Allocation." *Internat. J. Indus. Org.* 41:42–50.
- Blackwell, D. 1951. "Comparison of Experiments." In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, edited by J. Neyman, 93–102. Berkeley: Univ. California Press.
- Byford, M. C., and J. S. Gans. 2014. "Collusion at the Extensive Margin." *Internat. J. Indus. Org.* 37:75–83.
- Carlton, D. W., and J. M. Perloff. 1995. *Modern Industrial Organization*. 4th ed. Boston: Pearson.
- Chan, J., and W. Zhang. 2015. "Collusion Enforcement with Private Information and Private Monitoring." *J. Econ. Theory* 157:188–211.
- Compte, O. 1998. "Communication in Repeated Games with Imperfect Private Monitoring." *Econometrica* 66 (3): 597–626.
- EC (Commission of the European Communities). 2002a. "Commission Decision of 17/12/2002 Relating to a Proceedings under Article 81 of the EC Treaty and Article 53 of the EEA Agreement (Case COMP/E-2/37.667—Specialty Graphite)." http://ec.europa.eu/competition/antitrust/cases/dec_docs/37667/37667_86_1.pdf.
- . 2002b. "Commission Decision of 27.11.02 Relating to a Proceeding under Article 81 of the EC Treaty (Case COMP/E-1/37.152—Plasterboard)." http://ec.europa.eu/competition/antitrust/cases/dec_docs/37152/37152_72_1.pdf.
- . 2002c. "Commission Decision of 27 November 2002 Relating to a Proceeding Pursuant to Article 81 of the EC Treaty and Article 53 of the EEA Agreement (Case COMP/E-2/37.978—Methylglucamine)." *Official J. European Union* L 038:18–46. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32004D0104&from=EN>.
- . 2004. "Commission Decision of 3 September 2004 Relating to a Proceeding Pursuant to Article 81 of the EC Treaty and Article 53 of the EEA Agreement (Case COMP/E-1/38.069—Copper Plumbing Tubes)." http://ec.europa.eu/competition/antitrust/cases/dec_docs/38069/38069_121_4.pdf.
- Forges, F. 1986. "An Approach to Communication Equilibria." *Econometrica* 54 (6): 1375–85.
- Friedman, J. W. 1971. "A Non-cooperative Equilibrium for Supergames." *Rev. Econ. Studies* 38 (1): 1–12.
- Fuchs, W. 2007. "Contracting with Repeated Moral Hazard and Private Evaluations." *A.E.R.* 97 (4): 1432–48.
- Fudenberg, D., D. Levine, and E. Maskin. 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica* 62 (5): 997–1039.
- Genesove, D., and W. P. Mullin. 2001. "Rules, Communication, and Collusion: Narrative Evidence from the Sugar Institute Case." *A.E.R.* 91 (3): 379–98.
- Green, E. J. and R. H. Porter. 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52 (1): 87–100.
- Haltiwanger, J., and J. E. Harrington Jr. 1991. "The Impact of Cyclical Demand Movements on Collusive Behavior." *RAND J. Econ.* 22 (1): 89–106.
- Harrington, J. E., Jr. 2006. "How Do Cartels Operate?" *Found. and Trends Microeconomics* 2 (1): 1–105.
- Harrington, J. E., Jr., and J. Chen. 2006. "Cartel Pricing Dynamics with Cost Variability and Endogenous Buyer Detection." *Internat. J. Indus. Org.* 24 (6): 1185–1212.

- Harrington, J. E., Jr., and A. Skrzypacz. 2011. "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices." *A.E.R.* 101 (6): 2425–49.
- Hirshleifer, J. 1971. "The Private and Social Value of Information and the Reward to Inventive Activity." *A.E.R.* 61 (4): 561–74.
- Hochman, G., and E. Segev. 2010. "Managed Trade with Imperfect Information." *Internat. Econ. Rev.* 51 (1): 187–211.
- Hörner, J., and A. Jamison. 2007. "Collusion with (Almost) No Information." *RAND J. Econ.* 38 (3): 804–22.
- Kandori, M. 1991a. "Cooperation in Finitely Repeated Games with Imperfect Private Information." Working paper, Dept. Econ., Princeton Univ.
- . 1991b. "Correlated Demand Shocks and Price Wars during Booms." *Rev. Econ. Studies* 58 (1): 171–80.
- . 1992. "The Use of Information in Repeated Games with Imperfect Monitoring." *Rev. Econ. Studies* 59 (3): 581–93.
- Kandori, M., and H. Matsushima. 1998. "Private Observation, Communication and Collusion." *Econometrica* 66 (3): 627–52.
- Kim, D. 2016. "Comparison of Information Structures in Stochastic Games with Imperfect Public Monitoring." Working paper, Dept. Econ., Pennsylvania State Univ.
- Kloosterman, A. 2015. "Public Information in Markov Games." *J. Econ. Theory* 157:28–48.
- Kobayashi, H., and K. Ohta. 2012. "Optimal Collusion under Imperfect Monitoring in Multimarket Contact." *Games and Econ. Behavior* 76 (2): 636–47.
- Levenstein, M. C., and V. Y. Suslow. 2006. "What Determines Cartel Success?" *J. Econ. Literature* 44 (1): 43–95.
- Mailath, G. J., S. A. Matthews, and T. Sekiguchi. 2002. "Private Strategies in Finitely Repeated Games with Imperfect Public Monitoring." *Contributions Theoretical Econ.* 2 (1): 2. doi:10.2202/1534-5971.1046.
- Mailath, G. J., and L. Samuelson. 2006. *Repeated Games and Reputations: Long-Run Relationships*. Oxford: Oxford Univ. Press.
- Marshall, R. C., and L. M. Marx. 2012. *The Economics of Collusion*. Cambridge, MA: MIT Press.
- Matsushima, H. 2001. "Multimarket Contact, Imperfect Monitoring, and Implicit Collusion." *J. Econ. Theory* 98 (1): 158–78.
- Miyahara, Y., and T. Sekiguchi. 2013. "Finitely Repeated Games with Monitoring Options." *J. Econ. Theory* 148 (5): 1929–52.
- Myerson, R. B. 1986. "Multistage Games with Communication." *Econometrica* 54 (2): 323–58.
- Pesendorfer, M. 2000. "A Study of Collusion in First-Price Auctions." *Rev. Econ. Studies* 67 (3): 381–411.
- Phelan, C., and A. Skrzypacz. 2012. "Beliefs and Private Monitoring." *Rev. Econ. Studies* 79 (4): 1637–60.
- Rahman, D. 2014. "The Power of Communication." *A.E.R.* 104 (11): 3737–51.
- Rotemberg, J. J., and G. Saloner. 1986. "A Supergame-Theoretic Model of Price Wars during Booms." *A.E.R.* 76 (3): 390–407.
- Sekiguchi, T. 2002. "Existence of Nontrivial Equilibria in Repeated Games with Imperfect Private Monitoring." *Games and Econ. Behavior* 40 (2): 299–321.
- . 2015. "Multimarket Contact under Demand Fluctuations." *Internat. J. Game Theory* 44 (4): 1033–48.
- Selten, R. 1978. "The Chain Store Paradox." *Theory and Decision* 9 (2): 127–59.
- Spector, D. 2015. "Facilitating Collusion by Exchanging Non-verifiable Sales Reports." Working paper no. 2015-07, Paris School Econ.

- Stigler, G. J. 1964. "A Theory of Oligopoly." *J.P.E.* 72 (1): 44–61.
- Sugaya, T., and A. Wolitzky. 2017. "Bounding Equilibrium Payoffs in Repeated Games with Private Monitoring." *Theoretical Econ.* 12 (2): 691–729.
- . 2018a. "Bounding Payoffs in Repeated Games with Private Monitoring: n -Player Games." *J. Econ. Theory* 175:58–87.
- . 2018b. "The Revelation Principle in Multistage Games." Working paper, Dept. Econ., Massachusetts Inst. Tech.
- USDoJ/FTC (US Department of Justice/Federal Trade Commission). 2010. "Horizontal Merger Guidelines." <https://www.ftc.gov/sites/default/files/attachments/merger-review/100819hmg.pdf>.
- Whinston, M. D. 2006. *Lectures on Antitrust Economics*. Cambridge, MA: MIT Press.