

Collusion with Optimal Information Disclosure*

Takuo Sugaya Alexander Wolitzky

Stanford GSB MIT

January 28, 2026

Abstract

Motivated by recent concerns surrounding the use of third-party pricing algorithms by competing firms, we study repeated Bertrand competition where market demand or the cost of serving the market is observed by an intermediary (or “algorithm”) that selectively discloses demand or cost information to maximize firms’ collusive profit. We show that an *upper censorship* disclosure policy is optimal, which leads to *price rigidity* and *supra-monopoly prices* at some states. Improving the algorithm’s accuracy reduces expected consumer surplus whenever it does so under monopoly pricing. When the state is positively correlated over time, the algorithm discloses more information when recent demand was lower or costs were higher. The analysis extends to a generalized model that accommodates product differentiation and capacity constraints. We relate our findings to recent antitrust cases.

Keywords: collusion, information disclosure, pricing algorithms, price rigidity, consumer surplus

JEL codes: C73, D43, D82

*For helpful comments, we thank seminar participants at Bonn, ESSET, HKUST, JEA, SAET, SITE, VSET, and the 2025 Northwestern Antitrust Conference and MIT Sloan IO Theory Conference; and Kyle Bagwell, Alessandro Bonatti, Sophie Calder-Wang, Glenn Ellison, Drew Fudenberg, Joe Harrington, Anton Kolotilin, Leslie Marx, Jeanine Miklós-Thal, Stephen Morris, Volker Nocke, Juan Ortner, Jesse Shapiro, Jean Tirole, and Catherine Tucker.

1 Introduction

Firms increasingly use automated algorithms to set prices and other competitive variables, a development that has raised a range of regulatory and antitrust concerns (Mehra, 2015; Ezrachi and Stucke, 2017; Calvano et al., 2020a). A particular focus of some prominent recent cases is third-party algorithms that facilitate information-sharing among competing firms while recommending prices. For example, RealPage, Inc. is a company that markets revenue management software to commercial landlords. RealPage’s software gathers detailed, near real-time information on apartment prices and occupancy rates from its users and other non-public sources and uses this data—including data on market conditions gleaned from competitors—to recommend prices. Following a history of private litigation against RealPage, in August 2024 the US Department of Justice and eight state attorneys general sued RealPage, asserting that, “At bottom, RealPage is an algorithmic intermediary that collects, combines, and exploits landlords’ competitively sensitive information,” which constitutes an “unlawful scheme to decrease competition among landlords,” (USDOJ 2024a,b; see also Calder-Wang and Kim, 2024). Similar algorithmic intermediaries have emerged in a number of other industries, including retail gasoline pricing (A2i Systems and Kalibrate; see Assad et al., 2024) and hotel room pricing (IDeAS and Rainmaker; see Harrington, 2025). In addition, related concerns have also been raised regarding offline cartel facilitators, such as the Swiss consulting firm AC-Treuhand, which was prosecuted by the European Commission for facilitating several European industrial cartels by disclosing competitively sensitive information and recommending prices and market allocations (Harrington, 2006; Marshall and Marx, 2012).

Motivated by this type of setting, this paper develops a simple model of how an intermediary that possesses more detailed aggregate demand or cost information than its client firms can selectively disclose this information to maximize its clients’ collusive profit.¹ We work in the canonical setting of repeated Bertrand competition with stochastic demand, introduced

¹The model is intended as a benchmark and does not attempt to fully capture the complex industries mentioned above. For example, in practice the objective of an intermediary like RealPage may or may not be maximizing collusive profit, and the intermediary’s information may or may not be a superset of the firms’. We connect our analysis to the RealPage case in detail in Section 5.

by Rotemberg and Saloner (1986).² Following Rotemberg and Saloner, our baseline setting assumes homogeneous products and iid demand, although we subsequently relax both of these assumptions. To get a stark and tractable model, we assume that the current demand state is observed *only* by the intermediary (henceforth, the *algorithm*), which then discloses information about the state according to a known policy. We also make a key technical assumption that profit is affine in the unknown state, so that, for any fixed price and any distribution over states, expected profit is determined by the expected state. Under these assumptions, we characterize the disclosure policy and the (pure strategy, subgame perfect) equilibrium that maximizes the firms' profits.

The main result in our baseline model is that the (firm-)optimal policy is *upper censorship* together with *conditional monopoly pricing*: there is a cutoff demand state \hat{s} such that, if the current demand state s is below \hat{s} , the algorithm discloses s and recommends the corresponding monopoly price $p^m(s)$ to all firms; and if the current demand state s is above \hat{s} , the algorithm discloses only the event $\{s > \hat{s}\}$ and recommends the monopoly price conditional on this information, $p^m(\mathbb{E}[s|s > \hat{s}])$. The optimal equilibrium thus features *rigid prices*: prices are constant unless demand falls below \hat{s} . It also involves *supra-monopoly prices* for a range of demand states: for demand states s in the interval $(\hat{s}, \mathbb{E}[\tilde{s}|\tilde{s} > \hat{s}])$, the equilibrium price is $p^m(\mathbb{E}[\tilde{s}|\tilde{s} > \hat{s}])$, which is greater than the monopoly price in state s , $p^m(s)$, whenever the monopoly price $p^m(\cdot)$ is an increasing function of demand. Finally, as compared to the observable demand case studied by Rotemberg and Saloner, optimal collusive prices are higher—and consumer surplus is lower—for every demand state.

The logic of these results is as follows. As in Rotemberg and Saloner, firms are most tempted to undercut the collusive price when demand is high, as this is when the static monopoly profit $\Pi^m(s)$ is largest relative to the equilibrium continuation payoff. In Rotemberg and Saloner—which is identical to the special case of our model where the algorithm fully discloses the demand state—the cartel responds by reducing prices when demand is high, which reduces current-period profit and hence reduces the current-period deviation gain. (This is the logic of Rotemberg and Saloner's “price wars during booms.”) However,

²Stochastic demand and stochastic cost are equivalent up to a sign change in our model. For concreteness, we mostly discuss the stochastic demand case.

when an algorithm controls the firms' information, it is more profitable to reduce profit at high demand states by pooling these states with lower demand states and recommending the monopoly price conditional on the disclosed information, rather than cutting prices. In other words, the cartel reduces firms' temptation to deviate in high demand states *only* by reducing their information, not by reducing the recommended price conditional on their information. Technically, the key observation is that the firms' "capped monopoly profit," $\min \{\Pi^m(s), \Pi^{\max}\}$ —where Π^{\max} is the maximum industry profit that the firms can obtain in a single period without violating incentive constraints, which is independent of s with homogeneous products and iid demand—is a "convex-then-concave" function of s , and upper censorship is the optimal disclosure policy for a convex-then-concave objective function (Kolotilin, 2018; Dworczak and Martini, 2019; Kolotilin et al., 2022).

The optimal collusive equilibrium displays clean comparative statics. Reducing the number of firms, increasing the discount factor, or improving the algorithm's accuracy makes collusive prices more flexible and increases collusive profit. In addition, if improving the information of a monopoly seller reduces expected consumer surplus, then so does making collusive prices more flexible in our model. This result speaks directly to antitrust concerns regarding algorithmic information-sharing. Specifically, while prior studies have found an ambiguous effect of improved algorithmic demand prediction on consumer surplus (Sugaya and Wolitzky, 2018; Miklós-Thal and Tucker, 2019), our conclusion is more unambiguously negative. The reason is that prior studies assumed that the algorithm fully discloses its information to firms, while we assume that it selectively discloses its information to maximize firm profits, and therefore conceals information that would lead to price cuts if it were disclosed. Thus, while Miklós-Thal and Tucker (2019, p. 1553) find "somewhat reassuring results for antitrust authorities who are worried about the implications for anticompetitive and collusive behavior of the digital environment," we can unfortunately offer no such reassurances for algorithms that selectively disclose information to maximize collusive profit.

We relate our baseline model and results to some recent antitrust cases, focusing on the RealPage case. While the available evidence is limited, it seems consistent with our key mechanism wherein an intermediary with superior demand information selectively discloses this information to maximize its clients' joint profits, resulting in a stylized pricing pat-

tern over the business cycle where prices respond flexibly to demand conditions in industry downturns but are more rigid in upturns.

While our baseline model tractably generalizes Rotemberg and Saloner and delivers sharp and empirically plausible results, it does assume a rather special market structure. We therefore extend the model in three directions. First, we let the state persist over time, following a Markov process. Here the main results from the iid case go through, and there are also some new results. For example, we show that when demand is positively correlated over time, the algorithm discloses more information when recent demand was lower. (The opposite result holds with negative serial correlation.) The intuition is that with positive serial correlation, firms are more pessimistic about demand—and thus less tempted to deviate—when recent demand was lower, so the algorithm can disclose more information without prompting a deviation. We also show that the optimal collusive price is no longer always equal to the monopoly price for the disclosed mean demand, and that, while price is always monotone in current demand (as in the iid case and in contrast to Rotemberg and Saloner), it can be non-monotone in the previous period’s demand, so that the expected price conditional on the last-period demand can display countercyclicality similar to that in Rotemberg and Saloner.

Second, we consider a generalized model that accommodates product differentiation and capacity constraints.³ The insight that optimal information disclosure involves regions of censorship, price rigidity, and supra-monopoly pricing extends to the generalized model. However, the form of the optimal disclosure policy depends on details of the demand system and can differ from that in the baseline, homogeneous goods case. For example, with a symmetric, linear demand system, the optimal policy generally discloses the highest demand states as well as the lowest ones, while pooling a region of intermediate states. The intuition is that, with differentiated goods, the attainable profits at the highest demand states are high enough that it is optimal to disclose these states, even though the resulting prices must fall below the monopoly level to deter undercutting. Mathematically, the capped monopoly profit is now piecewise-convex rather than convex-then-concave, so in general an optimal

³These features are likely important in the RealPage case, although our discussion of this case in Section 5 focuses on our baseline model for simplicity.

disclosure policy censors an intermediate interval of states, rather than the highest states.

Third, we briefly consider the problem of designing a disclosure policy to maximize a weighted average of producer and consumer surplus, assuming that firms play their optimal equilibrium under the chosen policy. Here, we find that the consumer-optimal disclosure policy under linear demand with an unknown intercept is a binary signal that reveals only whether demand is below or above a cutoff, where the low signal induces the corresponding monopoly price, and the high signal sparks a “price war” (price below the corresponding monopoly price).

The remainder of the paper is organized as follows. Following a discussion of the literature, Section 2 presents the baseline model, Section 3 solves the model, and Section 4 discusses implications and comparative statics. Section 5 connects the analysis to recent antitrust cases, focusing on the RealPage case. Section 6 contains the extension to a persistent state and to a symmetric, linear demand system. Section 7 concludes and discusses further extensions. A generalized model that allows a range of demand systems as well as capacity constraints is presented in Appendix A. Finally, Appendix B considers more general objectives, such as maximizing consumer surplus.

Related literature. We contribute to the literatures on pricing algorithms, information-sharing among colluding firms, information design, and repeated games.

Much of the recent literature on pricing algorithms studies how independent algorithms can learn to set supra-competitive prices (Waltman and Kaymak, 2008; Calvano et al., 2020b; Klein, 2021; Asker, Fershtman, and Pakes, 2024; Banchio and Mantegazza, 2024), as well as the commitment value of adopting such algorithms (Cooper et al., 2015; Salcedo, 2015; Hansen, Misra, and Pai, 2021; Brown and MacKay, 2023; Lamba and Zhuk, 2024). We instead ask how a shared algorithm with demand information superior to the firms’ optimally discloses information to facilitate collusion. Sugaya and Wolitzky (2018, Example 3) and Miklós-Thal and Tucker (2019) show that the effect of disclosing demand information on collusive profit and consumer surplus is generally non-monotone, as it facilitates more accurate deviations as well as more accurate on-path pricing (a logic similar to Rotemberg and Saloner’s). O’Connor and Wilson (2019), Martin and Rasch (2022), and Bonatti, Fiocco,

and Piccolo (2024) document similar effects under imperfect monitoring.⁴ However, none of these papers characterizes optimal disclosure.

Harrington (2022) notes a reason why our model might *not* be a good fit for a third-party company like RealPage that designs and sells a pricing algorithm to competing firms: if firms independently decide whether to purchase and adopt the algorithm, a profit-maximizing algorithm designer's objective may be to maximize the *difference* in profit between adopters and non-adopters, rather than adopters' profits.⁵ This alternative objective could be considered in future research. Harrington (2025) considers a problem closer to ours, where the algorithm maximizes adopter's profit subject to the constraint that adoption is profitable. However, in his model, adopters commit to following the algorithm's price recommendations. Finally, a very recent paper by Harrington and Ortner (2025) considers a similar problem where adopters do not commit to following the recommendations. The key difference from our model is that Harrington and Ortner consider idiosyncratic demand uncertainty that washes out in aggregate. In that case (and assuming a linear demand system), they show that full information disclosure is optimal.

The broader literature on information-sharing among colluding firms considers a range of mechanisms, including the impact of improved monitoring (Abreu, Milgrom, and Pearce, 1991; Kandori, 1992; Harrington and Skrzypacz, 2011; Awaya and Krishna, 2016), the benefits of maintaining strategy uncertainty (Bernheim and Madsen, 2017; Sugaya and Wolitzky, 2018; Ortner, Sugaya, and Wolitzky, 2024; Kawai, Nakabayashi, and Ortner, 2024), and the allocative role of communication under incomplete information.⁶ These papers find that concealing various types of information can be advantageous for cartels. However, we are not aware of any prior work that studies optimal information disclosure for facilitating collusion.⁷

Optimal information disclosure has been studied extensively in static environments (Rayo

⁴Bonatti, Fiocco, and Piccolo (2024) focus on a comparison between revealing demand information before and after firms set prices.

⁵Bordoli (2025) and Wu (2025) consider a similar designer objective in a static setting.

⁶The latter literature contains papers on communicating private cost information (McAfee and McMillan, 1992; Athey and Bagwell, 2001; Athey, Bagwell, and Sanchirico, 2004; Skrzypacz and Hopenhayn, 2004), as well as private signals of stochastic market demand (Hanazono and Yang, 2007; Gerlach, 2009; Buehler and Gärtner, 2013; Sahueget and Walckiers, 2017).

⁷Hickok (2024) studies optimal information disclosure by a platform that takes a share of firms' revenue, finding that full disclosure is optimal.

and Segal, 2010; Kamenica and Gentzkow, 2011), especially in the affine case we focus on (Gentzkow and Kamenica, 2016; Kolotilin et al., 2017; Kolotilin, 2018; Dworczak and Martini, 2019), as well as in some specific dynamic settings (e.g., Ely, 2017; Renault, Solan, and Vieille, 2017). From a technical perspective, the closest paper is Kolotilin and Li (2021), who study a repeated cheap talk game with voluntary transfers. Like us, they reduce the problem of characterizing the optimal equilibrium to a static information design problem.⁸ However, this reduction works for different reasons in the two papers: in Kolotilin and Li, the key is the availability of transfers; for us, the key is the fact that static deviation gains are proportional to on-path payoffs under Bertrand competition.⁹ Kolotilin and Li's reduction also imposes a monotonicity constraint, which is absent in our setting. Kolotilin and Li (Proposition 4) derive conditions under which upper censorship is optimal, which include shape restrictions on utilities beyond affineness. For us, no conditions are required beyond affineness, due to the structure of Bertrand competition. Once the reduction to static information design (Lemma 1) is in place, our proof is essentially the same as Kolotilin and Li's (as well as Kolotilin, 2018; Dworczak and Martini, 2019; and others). However, the structure of Bertrand competition also lets us handle the Markov case in Section 6.1, and we also characterize more general optimal disclosure policies in Section 6.2 and Appendix A.

From the viewpoint of repeated game theory, we combine the recursive approach pioneered by Abreu (1986, 1988) with optimal public information disclosure. The approach is to consider static information design in the stage game augmented with continuation payoffs. This is especially tractable in symmetric games where symmetric equilibria can be shown to be without loss. The text of the paper illustrates this approach in two special settings (collusion with homogeneous goods or with differentiated goods with linear demand), while the model in Appendix A is a relatively general one where this methodology applies.

⁸Kuvalekar, Lipnowski, and Ramos (2022) also reduce a repeated communication game to a static one.

⁹In the generalized version of our model, the reduction works due to a more general one-to-one correspondence between static deviation gains and on-path payoffs.

2 A Model of Collusion with Information Disclosure

Prices and profits. We consider homogeneous-goods Bertrand competition among n firms with stochastic demand or a stochastic common production cost. In each period, a non-negative demand or cost state $s \in [\underline{s}, \bar{s}]$ is drawn independently from an atomless distribution F , and each firm i sets a non-negative price p_i , which is publicly observed. The firms' information about s is described below. The lowest-price firm i serves the entire market and makes profit $\Pi(p_i, s)$. The market is shared equally in case of a tie.

We focus on the case where s measures market demand. In this case, we assume that $\Pi(0, s) = 0$ for all s (normalizing costs to 0); that $\Pi(p, s)$ is continuous in p with a well-defined monopoly profit $\Pi^m(s) = \max_p \Pi(p, s)$ for each s ; and that $\Pi(p, s)$ is *affinely increasing* in s for each p : that is, $\Pi(p, \underline{s}) \leq \Pi(p, \bar{s})$ and

$$\Pi(p, s) = \frac{\bar{s} - s}{\bar{s} - \underline{s}} \Pi(p, \underline{s}) + \frac{s - \underline{s}}{\bar{s} - \underline{s}} \Pi(p, \bar{s}) \quad \text{for all } p, s.$$

In the alternative case where s measures a common production cost, the assumptions are slightly different. Here, we assume that $\Pi(p, s) \leq 0$ for all $p < s$ with equality at $p = s$; that $\Pi(p, s)$ is continuous in p with a well-defined monopoly profit $\Pi^m(s)$ for each s ; and that $\Pi(p, s)$ is affinely *decreasing* in s for each p . An interpretation of this case is that the firms are bidders in a procurement auction, where the cost of fulfilling the contract is privately observed by an intermediary who coordinates bid-rigging among the firms.¹⁰

Affineness in s is our key assumption. It has two important implications. First, expected profit is measurable with respect to mean demand: for any price p and any distribution of demand states $\mu \in \Delta([\underline{s}, \bar{s}])$, the expected profit from serving the market at price p is $\mathbb{E}^\mu[\Pi(p, s)] = \Pi(p, \mathbb{E}^\mu[s])$. Second, monopoly profit $\Pi^m(s) = \max_p \Pi(p, s)$ is increasing and convex in s as the maximum of increasing affine functions.¹¹

Affineness in s is a strong assumption, but it is satisfied in some important cases, which we return to throughout the paper. First, it holds if there is a binary underlying demand or cost state $\mathbf{s} \in \{\underline{s}, \bar{s}\}$, where s is a continuous signal of \mathbf{s} satisfying $\Pr(\mathbf{s} = \bar{s}|s) =$

¹⁰An example that fits this interpretation is the Kumatori Contractors Cooperative studied by Kawai, Nakabayashi, and Ortner (2024), which we discuss in Section 5.

¹¹In the stochastic cost case, $\Pi^m(s)$ is decreasing and convex in s .

$(s - \underline{s}) / (\bar{s} - \underline{s})$. Second, it holds for *linear demand with an unknown intercept*, where demand equals $D(p, s) = s - p$, and hence $\Pi(p, s) = p(s - p)$.¹² Third, it holds for *linear demand with an unknown constant marginal cost*, where demand equals $1 - p$ and marginal cost equals s , so that $\Pi(p, s) = (p - s)(1 - p)$.

Information. We assume that the firms do not directly observe the state s . Instead, s is observed by an intermediary—which we refer to as the *algorithm*—which maps s to a (possibly random) signal according to a known rule. We assume that the signal is publicly observed by all firms. Importantly, this assumption restricts the scope of our analysis to public information disclosure and rules out more general private communication.¹³ Since expected profit is measurable with respect to mean demand, it is without loss to view the algorithm as choosing a distribution G of the firms’ posterior expectations of s , which we denote by x . By Blackwell (1953) (see also Strassen, 1965; Kolotilin, 2018), such a distribution is consistent with Bayesian updating of the prior F if and only if $G \in MPC(F)$, the set of mean-preserving contractions of F . We refer to such a distribution G as a *disclosure policy*.

Repeated game equilibrium. The above game is repeated in discrete time with a common discount factor δ . In principle, the algorithm can choose a different disclosure policy G each period, but we will see that there is no benefit from doing so in the current model with an iid state.¹⁴

Our solution concept is pure strategy, subgame perfect equilibrium (henceforth, “equi-

¹²These two cases both nest Example 3 of Sugaya and Wolitzky (2018), which assumes a binary demand state and linear demand. The first case also nests the model of Miklós-Thal and Tucker (2019), which assumes a binary demand state and unit demand. Our analysis also applies for linear demand subject to a non-negativity constraint, $D(p, s) = \max\{s - p, 0\}$, so long as $\underline{s} \geq \bar{s}/2$, so that demand $D(p^m(s), s')$ is non-negative for any monopoly price $p^m(s)$ and demand state s' .

¹³With private signals, in each period the problem would become one of characterizing the optimal Bayes correlated equilibrium in a game with a continuum of states and actions and discontinuous payoffs. This problem is generally intractable. For example, see Smolin and Yamashita (2025) for results with concave payoffs, as well as a recent literature review. Private signals do improve on public ones—for example, Ortner, Sugaya, and Wolitzky (2024) find the optimal equilibrium with private signals in a similar model without demand uncertainty. Besides tractability, public communication and pure strategies ensure that the equilibrium will be robust to firms’ investigating their competitors’ prices before setting their own.

¹⁴In Section 6.1, the state follows a Markov process, and the optimal disclosure policy depends on the previous period’s state. In Section 6.2 and Appendix A, the optimal disclosure policy is time-invariant along the equilibrium path but discloses no information off path.

librium”). Here, pure strategies mean that, in each period, each firm i sets a deterministic price $p_i(x)$ as a function of the disclosed mean demand state x and the history of past mean demand states and all firms’ past prices.¹⁵

3 Optimal Information Disclosure and Pricing

We characterize the joint information disclosure and pricing policy that maximizes collusive profits (the sum of the firms’ payoffs). We reduce this problem to a static information design problem in Section 3.1 and solve it in Section 3.2.

3.1 Reduction to Static Information Design

For any number $V \geq 0$, define

$$\Pi^{\max}(\delta, n, V) = \frac{\delta V}{(1 - \delta)(n - 1)}. \quad (1)$$

Intuitively, this is the maximum profit Π such that a firm prefers to receive Π/n today and V/n in every future period rather than Π today and zero in the future. Thus, if the firms expect a future per-period collusive profit of V and the cartel tries to make profit $\Pi > \Pi^{\max}(\delta, n, V)$ in any period, a firm will deviate.

Next, define V^* as the greatest fixed point of the equation

$$V = \max_{G \in MPC(F)} \mathbb{E}^G [\min \{\Pi^m(x), \Pi^{\max}(\delta, n, V)\}]. \quad (2)$$

Intuitively, this is the maximum expected profit attainable by a joint information disclosure and pricing policy that never makes profit greater than $\Pi^m(x)$ (the maximum feasible profit at mean demand state x) or $\Pi^{\max}(\delta, n, V^*)$ (the maximum incentive compatible profit at mean demand state x , when future per-period collusive profit equals V^*). Note that the right-hand side of (2) is bounded by $\mathbb{E}^F [\Pi^m(s)]$, so V^* is well-defined.

¹⁵Restricting to pure strategies is standard but may not be without loss of generality, as randomization could deter deviations by making firms unsure of the winning price. This effect is studied in complete-information models by Bernheim and Madsen (2017) and Kawai, Nakabayashi, and Ortner (2024). Combining randomization and incomplete information is a possible direction for future research.

We show that optimal collusive profit equals V^* and that this profit level is attained by an equilibrium that is *symmetric* ($p_i(x)$ is always identical across firms i), *stationary* (the disclosure policy G is the same in every period and, on path, $p_i(x)$ is independent of the history of past demand realizations), and of a *grim trigger* form (play permanently reverts to the static Nash equilibrium following any deviation).

Lemma 1 *Optimal collusive profit equals V^* and is attained by a symmetric, stationary, grim trigger equilibrium. Moreover, a disclosure policy G is optimal if and only if it solves the maximization problem in (2) with $V = V^*$.*

Lemma 1 reduces the problem of finding an optimal equilibrium to the static information design problem on the right-hand side of equation (2), with $V = V^*$.

Proof. We first show that there exists a symmetric, stationary, grim trigger equilibrium that attains collusive profit V^* . For each x , let $p^m(x) \in \operatorname{argmax}_p \Pi(p, x)$ be a monopoly price in state x , and let

$$p(x) = \begin{cases} p^m(x) & \text{if } \Pi^m(x) \leq \Pi^{\max}(\delta, n, V^*), \\ \min \{p : \Pi(p, x) = \Pi^{\max}(\delta, n, V^*)\} & \text{if } \Pi^m(x) > \Pi^{\max}(\delta, n, V^*). \end{cases}$$

Note that $p(x)$ is well-defined by the intermediate value theorem, as $\Pi(0, x) = 0$ and $\Pi(p, x)$ is continuous in p .¹⁶ Let $G^* \in \operatorname{argmax}_{G \in MPC(F)} \mathbb{E}^G [\min \{\Pi^m(x), \Pi^{\max}(\delta, n, V^*)\}]$. Consider disclosure policy G^* , together with the strategy profile where all firms price at $p(x)$ whenever mean demand x realizes on path, and all firms price at zero off path. This is a symmetric, stationary, grim trigger strategy profile, which yields collusive profit V^* by construction. To see that it is an equilibrium, note that a firm's best deviation when realized mean demand is x is to price just below $p(x)$: this is immediate if $p(x) = p^m(x)$ and otherwise follows because $p(x)$ is the smallest price p satisfying $\Pi(p, x) = \Pi^{\max}(\delta, n, V^*)$, so that $\Pi(p', x) < \Pi^{\max}(\delta, n, V^*)$ for all $p' < p(x)$. This deviation wins the entire market in the current period, but forfeits an expected profit of V^*/n in every future period. Thus, the strategy profile is

¹⁶Here and throughout, we write proofs for the case where s is a demand state. The proofs for the case where s is a cost state are nearly identical.

an equilibrium if and only if, for all x , we have

$$\begin{aligned} (1 - \delta) \Pi(p(x), x) + \delta(0) &\leq \frac{1}{n} ((1 - \delta) \Pi(p(x), x) + \delta V^*) \iff \\ \Pi(p(x), x) &\leq \frac{\delta V^*}{(1 - \delta)(n - 1)} = \Pi^{\max}(\delta, n, V^*). \end{aligned}$$

Since this inequality holds by construction, the strategy profile is an equilibrium.

We now show that no equilibrium yields higher profit. Fix any equilibrium, and let \bar{V} be the supremum over periods t and histories of play up to and including period t of the expected per-period collusive profit from period $t + 1$ onward. Now fix an arbitrary period t and a history of play up to period t , and suppose that when the realized mean demand in period t at this history is x , the prescribed winning price is $p(x)$, and each firm i wins with probability α_i and obtains equilibrium continuation value v_i . (So, $\alpha_i = 1/|j : p_j(x) = p(x)|$ if $p_i(x) = p(x)$, and $\alpha_i = 0$ otherwise. Note that each $p_i(x)$ —and thus the winning price $p(x)$ —is deterministic by our restriction to pure strategy equilibria.) Since a possible deviation for firm i is to price just below $p(x)$ and a firm's minimax payoff is zero, firm i 's incentive constraint implies

$$(1 - \delta) \Pi(p(x), x) + \delta(0) \leq \alpha_i (1 - \delta) \Pi(p(x), x) + \delta v_i.$$

Averaging this inequality over the n firms, we have

$$(1 - \delta) \Pi(p(x), x) \leq \frac{1}{n} \left((1 - \delta) \Pi(p(x), x) + \delta \sum_i v_i \right) \leq \frac{1}{n} ((1 - \delta) \Pi(p(x), x) + \delta \bar{V}),$$

where the second inequality is by definition of \bar{V} . This inequality is equivalent to $\Pi(p(x), x) \leq \Pi^{\max}(\delta, n, \bar{V})$. Since we also have $\Pi(p(x), x) \leq \Pi^m(x)$ by definition, and these inequalities hold for any x , expected collusive profit in period t is at most $\max_{G \in MPC(F)} \mathbb{E}^G [\min \{\Pi^m(x), \Pi^{\max}(\delta, n, \bar{V})\}]$. Since this holds for any period t , we have $\bar{V} \leq \max_{G \in MPC(F)} \mathbb{E}^G [\min \{\Pi^m(x), \Pi^{\max}(\delta, n, \bar{V})\}]$. But this implies that equation (2) has a fixed point weakly above \bar{V} , and hence $\bar{V} \leq V^*$, by definition of V^* . ■

A direct implication of Lemma 1 is that collusion is impossible if $\delta < (n - 1)/n$. (The

same condition implies that collusion is impossible under full information disclosure, as in Rotemberg and Saloner.) Conversely, if $\delta \geq (n-1)/n$ then monopoly profit under no information disclosure, $\Pi^m(\mathbb{E}^F[s])$, is attainable.

Lemma 2 *If $\delta < (n-1)/n$ then $V^* = 0$. Conversely, if $\delta \geq (n-1)/n$ then $V^* \geq \Pi^m(\mathbb{E}^F[s])$.*

Proof. If $\delta < (n-1)/n$ then $\Pi^{\max}(\delta, n, V) < V$ for all $V > 0$, so the only solution to (2) is $V = 0$. Conversely, if $\delta \geq (n-1)/n$ then no information disclosure together with a constant on-path price of $p^m(\mathbb{E}^F[s])$ and zero prices off path is an equilibrium, with expected profit $\Pi^m(\mathbb{E}^F[s])$. ■

Given Lemma 2, we henceforth assume that $\delta \geq (n-1)/n$.

3.2 Optimality of Upper Censorship

The information design problem in (2) is easily solved using recent results from the static information design literature.

First, let $s^* \in [s, \bar{s}]$ solve

$$\Pi^m(s^*) = \Pi^{\max}(\delta, n, V^*) \quad (3)$$

if such a demand state exists, and let $s^* = \bar{s}$ otherwise. Note that, by Lemma 2 and our assumption that $\delta \geq (n-1)/n$, we have $V^* \geq \Pi^m(\mathbb{E}^F[s])$, and hence $s^* \geq \mathbb{E}^F[s]$.¹⁷ Thus, since F is atomless, there exists $\hat{s} \in [s, s^*]$ such that

$$\mathbb{E}^F[s | s \geq \hat{s}] = s^*.$$

We can now characterize the optimal information disclosure and pricing policy. Figure 1 illustrates the optimal disclosure policy, as well as the construction of s^* and \hat{s} .

Theorem 1 *With stochastic demand, the unique optimal disclosure policy is the upper censorship policy that discloses demand states below \hat{s} and conceals demand states above \hat{s} . The*

¹⁷This follows because if $s^* < \mathbb{E}^F[s]$ then $\Pi^m(\mathbb{E}^F[s]) > \Pi^{\max}(\delta, n, V^*)$ by (3) and monotonicity of $\Pi^m(s)$, but then we would have $V^* > \Pi^{\max}(\delta, n, V^*)$ by Lemma 2, contradicting the definition of V^* .

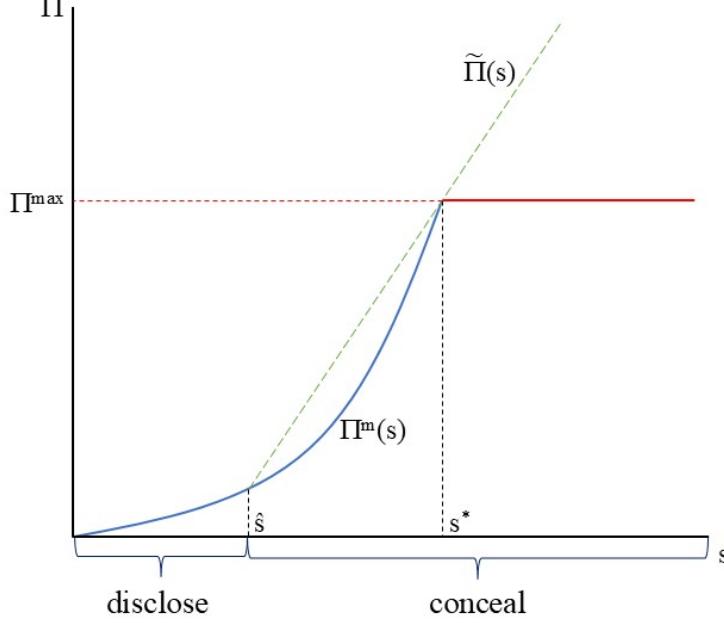


Figure 1: The optimal disclosure policy with undifferentiated products. First, s^* is determined as the solution to $\Pi^m(s^*) = \Pi^{\max}(\delta, n, V^*)$. Then, \hat{s} is determined as the solution to $E^F[s|s \geq \hat{s}] = s^*$. The optimal policy discloses demand states $s < \hat{s}$ and recommends the corresponding monopoly price, $p^m(s)$; and conceals demand states $s \geq \hat{s}$ and recommends the monopoly price conditional on this information, $p^m(s^*)$. The function $\tilde{\Pi}(s)$ is defined in the proof of Theorem 1.

unique optimal collusive price $p(s)$ in state s is given by

$$p(s) = \begin{cases} p^m(s) & \text{if } s < \hat{s}, \\ p^m(s^*) & \text{if } s \geq \hat{s}. \end{cases} \quad (4)$$

With stochastic costs, the unique optimal disclosure policy is the analogous lower censorship policy that discloses cost states above \hat{s} satisfying $\mathbb{E}^F[s|s \leq \hat{s}] = s^*$ and conceals cost states below \hat{s} .¹⁸ The unique optimal collusive price $p(s)$ in state s is given by (4) with the reversed inequalities.

With stochastic demand, note that $\hat{s} = \underline{s}$ —so no disclosure is optimal—iff $s^* = \mathbb{E}^F[s]$, which holds iff $\delta = (n - 1)/n$. Conversely, $\hat{s} = \bar{s}$ —so full disclosure is optimal—iff $s^* = \bar{s}$, which holds iff $\Pi^m(\bar{s}) \leq \Pi^{\max}(\delta, n, V^*)$. Otherwise, we have $\Pi^m(\mathbb{E}^F[s]) < \Pi^{\max}(\delta, n, V^*) < \Pi^m(\bar{s})$, and partial disclosure is optimal.

To understand Theorem 1, note that disclosing demand information increases expected

¹⁸The definitions of Π^{\max} and s^* remain as in (1) and (3).

monopoly profits—as $\Pi^m(s)$ is convex—but revealing that expected demand is too high requires cutting price to deter a deviation (as in Rotemberg and Saloner). The theorem says that it is optimal to disclose low demand states and conceal high ones, such that the mean concealed state s^* is the highest state s that does not require a price cut from the corresponding monopoly price $p^m(s)$ to deter a deviation.

The intuition is that it cannot be optimal to disclose a mean demand state $s > s^*$, as pooling s with lower demand states would increase expected profit. In particular, full disclosure together with price cuts during booms—Rotemberg and Saloner’s equilibrium—is suboptimal in our model. In addition, it cannot be optimal to pool two demand states below s^* , as separating these states increases expected profit since $\Pi^m(s)$ is convex. Finally, it is more profitable to pool demand states above s^* with intermediate states $s \in [\hat{s}, s^*]$ rather than low states $s < \hat{s}$, as this spreads out the distribution of disclosed mean demand states s on the interval $[\underline{s}, s^*]$, where $\Pi^m(s)$ is convex.

A more technical explanation is that the objective function $\min \{\Pi^m(s), \Pi^{\max}(\delta, n, V^*)\}$ is increasing and convex in s for $s \leq s^*$ and is constant in s for $s > s^*$. Thus, (2) describes a mean-measurable information design problem with an objective function that is “S-shaped”: first convex, then concave. It is well-known that the solution to such a problem is upper censorship (e.g., Kolotilin, 2018; Dworczak and Martini, 2019; Kolotilin and Li, 2021; Kolotilin et al., 2022). Moreover, adapting the standard proofs to the current setting where the objective function is not only convex-then-concave but convex-then-constant implies that the solution must take the prescribed form, where the mean censored state s^* lies at the kink of the objective function.¹⁹ In particular, the disclosed mean demand state s always satisfies $\Pi^m(s) \leq \Pi^{\max}(\delta, n, V^*)$, which implies that conditional monopoly pricing is optimal.

Proof. Define an auxiliary objective function

$$\tilde{\Pi}(s) = \begin{cases} \Pi^m(s) & \text{if } s < \hat{s}, \\ \frac{s^* - s}{s^* - \hat{s}} \Pi^m(\hat{s}) + \frac{s - \hat{s}}{s^* - \hat{s}} \Pi^m(s^*) & \text{if } s \geq \hat{s}. \end{cases}$$

Note that $\tilde{\Pi}(s)$ is convex and $\tilde{\Pi}(s) \geq \min \{\Pi^m(s), \Pi^{\max}(\delta, n, V^*)\}$ for all s . (See Figure 1.)

¹⁹Specifically, we adapt the proofs of Proposition 3 of Dworczak and Martini (2019) and Proposition 4 of Kolotilin and Li (2021). Theorem 1 is also a special case of Theorem 4 in Appendix A.

Consider the auxiliary problem, $\max_{G \in MPC(F)} \mathbb{E}^G \left[\tilde{\Pi}(s) \right]$. Since $\tilde{\Pi}(s)$ is convex, the solution is full disclosure ($G = F$), and the resulting value is

$$\begin{aligned} \mathbb{E}^F \left[\tilde{\Pi}(s) \right] &= F(\hat{s}) \mathbb{E}^F [\Pi^m(s) | s < \hat{s}] \\ &\quad + (1 - F(\hat{s})) \mathbb{E}^F \left[\frac{s^* - s}{s^* - \hat{s}} \Pi^m(\hat{s}) + \frac{s - \hat{s}}{s^* - \hat{s}} \Pi^m(s^*) | s \geq \hat{s} \right] \\ &= F(\hat{s}) \mathbb{E}^F [\Pi^m(s) | s < \hat{s}] + (1 - F(\hat{s})) \Pi^m(s^*). \end{aligned}$$

Since $\tilde{\Pi}(s) \geq \min \{ \Pi^m(s), \Pi^{\max}(\delta, n, V^*) \}$ for all s , this value is an upper bound for $\max_{G \in MPC(F)} \mathbb{E}^G [\min \{ \Pi^m(s), \Pi^{\max}(\delta, n, V^*) \}]$. But it is attained by upper censorship with cutoff \hat{s} , so this policy is optimal. Moreover, this policy is the unique one that induces only posteriors s where $\tilde{\Pi}(s) = \min \{ \Pi^m(s), \Pi^{\max}(\delta, n, V^*) \}$, so it is the unique optimal policy. Finally, this disclosure policy is optimal only in conjunction with the prescribed prices. ■

4 Implications and Comparative Statics

We now discuss the baseline model's predictions, comparative statics, and consumer welfare implications.

4.1 Model Predictions and Impact of Selective Disclosure

The optimal disclosure and pricing policy characterized in Theorem 1 has the following features. In the subsequent discussion, we assume that the monopoly price $p^m(s)$ is unique and increasing in s .

Conditional monopoly pricing. Along the equilibrium path of play, the cartel prevents deviations solely by reducing its members' information, not by reducing collusive prices below the monopoly price conditional on their information. Consequently, collusive prices for a cartel aided by an algorithm that observes a state with distribution F is identical to monopoly pricing for a single firm that observes a state with a less informative distribution G^* . This observation will be useful for the comparative statics and consumer welfare results in Section 4.2, as it implies that comparative statics for collusive prices are equivalent to

comparative statics for monopoly prices with respect to the monopoly’s information.

Price rigidity—not price wars—during booms. The collusive price $p(s)$ is increasing for $s < \hat{s}$ and is constant (at a higher level) for $s \geq \hat{s}$. In particular, collusive prices exhibit rigidity for demand states $s \geq \hat{s}$, rather than “price wars” as in Rotemberg and Saloner. This result gives a novel rationale for oligopoly price rigidity in intermediated collusion settings: prices are rigid because colluding firms optimally limit their own information about market conditions to deter deviations.²⁰

In addition, this result shows that algorithmic pricing does not unambiguously increase price flexibility under collusion. This result somewhat goes against the received view following Chen, Mislove, and Wilson’s (2016) study of retailer pricing on Amazon, which finds that very flexible prices are a hallmark of algorithmic pricing under competition.

Supra-monopoly pricing. Collusive prices are *above monopoly* at intermediate demand states: for $s \in (\hat{s}, s^*)$, the optimal collusive price is $p(s) = p^m(s^*) > p^m(s)$. Moreover, these demand states satisfy $\Pi^m(s) < \Pi^{\max}(\delta, n, V^*)$, so monopoly profit can be attained at any one of these states s by disclosing s and recommending price $p^m(s)$ (holding the rest of the equilibrium fixed). Thus, for a range of demand states where monopoly profit is attainable, the algorithm instead implements supra-monopoly prices that deliver lower profits. The reason why is that recommending the supra-monopoly price $p^m(s^*) > p^m(s)$ in states $s \in (\hat{s}, s^*)$ lets the algorithm recommend the same price in states $s > s^*$, where this price would be too high to be incentive compatible if the state were disclosed. In other words, price rigidity for demand states above \hat{s} results in an inefficiently high price for demand states in (\hat{s}, s^*) , but thereby supports a higher price for demand states above s^* than would be attainable under full information.²¹

Impact of selective information disclosure. The algorithm’s ability to conceal information

²⁰While not focused on intermediated cartels per se, Carlton (1986) and others find that prices are more rigid in concentrated industries, and Harrington (2008) and others suggest price rigidity as a collusive marker. Existing theories of rigid collusive prices include Athey, Bagwell, and Sanchirico (2004) and Hanazono and Yang (2007) (based on incentive costs of inducing firms to reveal private cost or demand information) and Maskin and Tirole (2001) (who model “kinked demand curves” as a result of Markov perfect equilibria with staggered price setting).

²¹Supra-monopoly pricing at intermediate demand states is analogous to “over-pooling”—where first-best actions are not taken even in some states where they are implementable—in Kolotilin and Li (2021).

leads to higher prices and lower consumer surplus in every demand state and thus unambiguously harms consumers. To see this, let V^{FD} be optimal collusive profit under full disclosure, which is given by the greatest fixed point of the equation

$$V^{FD} = \mathbb{E}^F \left[\min \left\{ \Pi^m(s), \Pi^{\max}(\delta, n, V^{FD}) \right\} \right],$$

and let s^{FD} solve

$$\Pi^m(s^{FD}) = \Pi^{\max}(\delta, n, V^{FD}).$$

As in Rotemberg and Saloner, optimal collusive prices under full disclosure are given by

$$p^{FD}(s) = \begin{cases} p^m(s) & \text{if } s \leq s^{FD}, \\ \min \{p : \Pi(p, s) = \Pi^{\max}(\delta, n, V^{FD})\} & \text{if } s > s^{FD}. \end{cases}$$

Since $p^m(s)$ is increasing and $\Pi(p, s)$ is increasing in s , it follows that $p^{FD}(s)$ is increasing for $s < s^{FD}$ and decreasing for $s \geq s^{FD}$. The latter “price wars during booms” result is Rotemberg and Saloner’s key message.

Proposition 1 *As compared to collusive prices under full disclosure, collusive prices under the optimal disclosure policy are higher at each demand state. Consequently, selective information disclosure reduces consumer surplus.*

Proof. Note that $V^{FD} \leq V^*$, and hence $s^{FD} \leq s^*$. Therefore, letting $p(s)$ be the optimal collusive price in (4), for $s < s^*$ we have $p(s) \geq p^m(s) \geq p^{FD}(s)$, and for $s > s^*$ we have $p(s) = p^m(s^*) \geq p^m(s^{FD}) \geq p^{FD}(s)$, where the first inequality holds because $s^{FD} \leq s^*$ and p^m is increasing, and the second holds because $s^{FD} \leq s$ and p^{FD} is decreasing for $s \geq s^{FD}$.

■

Empirical predictions, collusive markers, and the interpretation of price wars. The baseline model has three main empirical predictions:

1. The support of the distribution of equilibrium prices consists of an interval $[p^m(\underline{s}), p^m(\hat{s})]$ and a single higher price $p^m(s^*)$.

2. Prices are rigidly fixed at $p^m(s^*)$ for all demand states except the lowest ones. For low demand states, prices are discretely lower than $p^m(s^*)$ but vary flexibly in the interval $[p^m(\underline{s}), p^m(\hat{s})]$. Overall, prices are pro-cyclical: $p(s)$ is non-decreasing.
3. While prices are pro-cyclical, the gap between monopoly and collusive prices, $p^m(s) - p(s)$, is non-monotone: first zero, then negative, then positive.

The predicted form of price rigidity—a rigid, high price together with an interval of flexible lower prices—is distinctive to our model and is thus a possible collusive marker.

The pro-cyclical relationship between prices and demand in our model gives an alternative interpretation of the “price wars” predicted by Green and Porter (1984) and other models of collusion under imperfect monitoring. In Green and Porter, prices are pro-cyclical: prices drop following demand shortfalls as part of an optimal repeated game equilibrium under imperfect monitoring. In contrast, in Rotemberg and Saloner, prices are counter-cyclical in high demand states. Interestingly, while our model is much closer to Rotemberg and Saloner’s, our prediction of pro-cyclical prices coincides with Green and Porter’s, albeit by a different mechanism: perfect monitoring and selectively disclosed demand information, rather than imperfect monitoring. This observation is relevant for a line of papers that have tested the competing predictions of Green and Porter and Rotemberg and Saloner (e.g., Porter, 1983; Ellison, 1994) and have typically found results more favorable to Green and Porter’s prediction of pro-cyclical prices. Relative to this literature, our analysis shows that perfect monitoring and selective information disclosure is an alternative explanation for pro-cyclical prices.

One way to distinguish our theory from Green and Porter’s would be to estimate the gap between monopoly and collusive prices, $p^m(s) - p(s)$, over the cycle. In Green and Porter, the gap is larger following low demand states: collusion is “more successful” when recent demand was high. In our model, the gap is larger in high demand states (and can even be negative): collusion is more successful when demand is low.²² It would be interesting to test

²²With the exception of the distinctive prediction of supra-monopoly prices at intermediate demand states, our prediction that the gap $p^m(s) - p(s)$ is pro-cyclical is as in Rotemberg and Saloner. Thus, our model and Rotemberg and Saloner’s make different predictions about prices, $p(s)$ (non-decreasing in our model; single-peaked in theirs), but similar predictions about the difference $p^m(s) - p(s)$ (non-decreasing in both models, with the exception of an intermediate region of supra-monopoly pricing in ours).

these predictions.

4.2 Comparative Statics and Consumer Welfare

We now turn to comparative statics. In what follows, we say that collusive prices are *more flexible* if the optimal disclosure policy G^* spreads out in the mean-preserving spread sense. By Blackwell (1953), this is equivalent to the algorithm's output becoming more informative. Holding fixed the algorithm's information F , this is also equivalent to increasing the censorship cutoff \hat{s} , so a wider range of demand states are disclosed.²³ We also continue to assume that $\delta \geq (n - 1)/n$, as otherwise collusion is impossible by Lemma 2.

Proposition 2 *Collusive profit V^* is higher and collusive prices are more flexible when*

1. *there are fewer firms (n decreases),*
2. *the firms are more patient (δ increases), or*
3. *the algorithm is more accurate (F increases in the mean-preserving spread sense).*

The intuition for the first two results is that decreasing n or increasing δ relaxes the firms' incentive constraints, which lets the algorithm disclose a wider range of states without prompting a deviation. In addition, a more accurate algorithm generates greater collusive profits, which again relaxes incentive constraints and allows greater disclosure.²⁴

Proof. For the first two results, note that n and δ affect V^* and p only through the function Π^{\max} , which is decreasing in n and increasing in δ . Thus, decreasing n or increasing δ shifts up the right-hand side of (2) as a function of V , which increases the greatest fixed point V^* . In turn, an increase in V^* increases s^* and \hat{s} , which makes prices more flexible.

For the third result, we prove in the appendix that a more informative prior implies a more informative optimal disclosure policy in static information design problems with a convex-constant objective: for any distributions (F_1, F_2, G_1, G_2) where F_2 is a mean-preserving

²³If F changes (as in Proposition 2.3, the comparative static with respect to the algorithm's information), then G^* can spread out even as \hat{s} decreases.

²⁴A similar effect is documented in Theorem 5 of Harrington (2025). Improving the algorithm's information corresponds to taking a mean-preserving spread of F by Blackwell (1953)—the interpretation is that there is a distribution H of an underlying state \mathbf{s} , and the distribution of the algorithm's expectation of the underlying state is F , a mean-preserving contraction of H .

spread of F_1 , G_1 is the distribution of x under an optimal disclosure policy for prior F_1 , and G_2 is the distribution of x under an optimal disclosure policy for prior F_2 , we have that G_2 is a mean-preserving spread of G_1 . Given this result, spreading out F implies that the optimal disclosure policy on the right-hand side of (2) is more informative for any fixed value for V . In addition, since $MPC(F_1) \subseteq MPC(F_2)$, spreading out F clearly shifts up the right-hand side of (2) as a function of V , which increases the greatest fixed point V^* . Since increasing V also implies a more informative optimal disclosure policy for any fixed prior (as is obvious from Theorem 1), the result follows. ■

Proposition 2 gives sufficient conditions for collusive prices to become more flexible. In general, the effect of price flexibility on expected consumer surplus is ambiguous. However, since collusive prices equal monopoly prices for a monopoly facing state distribution G^* , price flexibility benefits consumers if and only if improving the information of a monopoly does the same.

To state this result, we must assume that expected consumer surplus is measurable with respect to the distribution of posterior mean states x . To this end, let $CS(p, s)$ denote consumer surplus at price p and state s . We say that consumer surplus is *quasi-linear* in s if there exist functions f , g , and h such that

$$CS(p, s) = f(s) + g(p)s + h(p) \quad \text{for all } p, s.$$

For example, under linear demand with an unknown intercept, we have $CS(p, s) = (s - p)^2/2 = f(s) + g(p)s + h(p)$, where $f(s) = s^2/2$, $g(p) = -p$, and $h(p) = p^2/2$; and under linear demand with an unknown constant marginal cost, we have $CS(p, s) = (1 - p)^2/2 = f(s) + g(p)s + h(p)$, where $f(s) = 0$, $g(p) = 0$, and $h(p) = (1 - p)^2/2$.

Quasi-linearity implies that expected consumer surplus is measurable with respect to the distribution of posterior mean states x , whenever prices $p(x)$ are measurable with respect to x . To see this, take any distribution $\tau \in \Delta(\Delta([\underline{s}, \bar{s}]))$ of distributions $\mu \in \Delta([\underline{s}, \bar{s}])$ of the state \tilde{s} such that $\mathbb{E}^\tau[\mu] = F$. Letting $x_\mu = \mathbb{E}^\mu[s]$ be the mean state under distribution

μ and letting G be the distribution of mean state x , we have

$$\begin{aligned}\mathbb{E}^\tau [\mathbb{E}^\mu [CS(p(x_\mu), s)]] &= \mathbb{E}^\tau [\mathbb{E}^\mu [f(s) + g(p(x_\mu))s + h(p(x_\mu))]] \\ &= \mathbb{E}^F [f(s)] + \mathbb{E}^\tau [g(p(x_\mu))x_\mu + h(p(x_\mu))] \\ &= \mathbb{E}^F [f(s)] + \mathbb{E}^G [g(p(x))x + h(p(x))],\end{aligned}$$

where the second line uses the law of iterated expectation. Thus, the expected consumer surplus under any disclosure policy and (mean-measurable) pricing policy equals $\mathbb{E}^G [g(p(x))x + h(p(x))]$, where G is the distribution of the disclosed mean state x , plus a constant $\mathbb{E}^F [f(s)]$.

In particular, if $g(p^m(x))x + h(p^m(x))$ is concave (resp., convex) in x , then expected consumer surplus under monopoly pricing is higher when the monopoly has less (resp., more) information. Thus, since collusive prices when the algorithm observes a state with distribution F equal monopoly prices when the monopoly observes a state with distribution G^* , the concavity or convexity of $g(p^m(x))x + h(p^m(x))$ determines the implications of the comparative statics in Proposition 2 for expected consumer surplus (with the caveat that increasing n or decreasing δ always benefits consumers if it causes δ to fall below $(n - 1)/n$, rendering collusion impossible).

Proposition 3 *Assume that consumer surplus is quasi-linear in x . If $g(p^m(x))x + h(p^m(x))$ is concave (resp., convex) in x , then expected consumer surplus is lower (resp., higher) when there are fewer firms, the firms are more patient, or the algorithm is more accurate (resp., so long as $\delta \geq (n - 1)/n$).*

Proof. Immediate from Proposition 2 and Jensen's inequality. ■

As a corollary, we obtain the corresponding consumer surplus implications under linear demand with an unknown intercept or unknown constant marginal cost.

Corollary 1 *Under linear demand with an unknown intercept, expected consumer surplus is lower when there are fewer firms, the firms are more patient, or the algorithm is more accurate.*

Under linear demand with an unknown constant marginal cost, expected consumer surplus is higher when there are fewer firms, the firms are more patient, or the algorithm is more accurate.

accurate (so long as $\delta \geq (n - 1)/n$).

Proof. Under linear demand with an unknown intercept, we have $p^m(x) = x/2$ and $CS(p, s) = f(s) + g(p)s + h(p)$, where $f(s) = s^2/2$, $g(p) = -p$, and $h(p) = p^2/2$. Hence,

$$g(p^m(x))x + h(p^m(x)) = -\frac{x^2}{2} + \frac{x^2}{8} = -\frac{3x^2}{8},$$

a concave function of x .

Under linear demand with an unknown constant marginal cost, we have $p^m(x) = (1 + x)/2$ and $CS(p, s) = f(s) + g(p)s + h(p)$, where $f(s) = 0$, $g(p) = 0$, and $h(p) = (1 - p)^2/2$. Hence,

$$g(p^m(x))x + h(p^m(x)) = \frac{(1 - x)^2}{8},$$

a convex function of x . ■

Proposition 3 shows that improving the algorithm's accuracy reduces consumer surplus whenever improving a monopoly's information does so, which Corollary 1 shows holds under linear demand with an unknown intercept.²⁵ This finding contrasts with results of Sugaya and Wolitzky (2018, Example 3) and Miklós-Thal and Tucker (2019), who find that a more accurate demand prediction algorithm increases consumer surplus when the firms' discount factor δ lies in an intermediate range. The reason for the difference is that those papers assume that the algorithm fully discloses its information, which sparks price wars during booms. In contrast, with optimal information disclosure, a more accurate algorithm makes prices more flexible without triggering price wars, which reduces expected consumer surplus whenever improving a monopoly's information does so. Our assessment of the likely impact of improved algorithmic demand prediction on consumer surplus is thus considerably more pessimistic.

However, improving the algorithm's accuracy also *increases* consumer surplus whenever improving a monopoly's information does so, which Corollary 1 shows holds under linear

²⁵Versions of the result that improving a monopoly's information about the intercept of a linear demand curve decreases expected consumer surplus were shown by Pigou (1920), Vives (2001), and Farboodi, Haghpanah, and Shourideh (2025).

demand with an unknown constant marginal cost.²⁶ In this case, expected consumer surplus also increases when there are fewer or more patient firms—so long as $\delta \geq (n - 1)/n$, so that collusion on the no-disclosure monopoly price $p^m(\mathbb{E}^F[s])$ is an equilibrium. The explanation is that, once $\delta \geq (n - 1)/n$, collusive prices equal monopoly prices for some information structure, so when consumers benefit from improving a monopoly’s information, they also benefit from the more flexible collusive prices that result from reducing n or increasing δ .

5 Connections to Recent Algorithmic Pricing Cases

Our model and results can be connected to recent algorithmic pricing cases. The cases are of course more complicated than any model, but we argue that the mechanisms at the core of our model appear appropriate to these settings and our predictions seem broadly consistent with the market outcomes.

Our model is motivated by and most directly relates to third-party algorithms that facilitate information-sharing among competitors and recommend prices, like RealPage in apartment rentals, IDeaS and Rainmaker in hotel rooms, and A2i Systems and Kalibrate in retail gasoline. Of these cases, the most information is available for RealPage. We therefore focus on this case, discussing some others briefly at the end of this section.

We focus on three aspects of the RealPage case. First, we argue that this case fits most key assumptions of our model setup. (The most important exception in our view is that algorithmic pricing adoption is far from universal among apartment management companies, and our model abstracts from competition with non-adopters as well as firms’ decision to adopt algorithmic pricing.²⁷) Second, we argue that RealPage intentionally withholds some information from its clients, and that this is plausibly done to make it easier to control its clients’ pricing (and specifically to prevent them from undercutting RealPage’s recommended prices). Third, we argue that the resulting price dynamics over the business cycle appear consistent with our prediction of price cuts during downturns and price rigidity during booms.

²⁶A version of the result that improving a monopoly’s information about its constant marginal cost increases expected consumer surplus was shown by Vives (2001).

²⁷These features are a key focus of the complementary work of Harrington (2025) and Harrington and Ortner (2025). In turn, these papers abstract from firms’ incentives to set their recommended prices and from aggregate demand uncertainty, respectively, which are key aspects of our model.

To make these points, we draw primarily on the 2024 US DOJ complaint and the subsequent 2025 proposed settlement (USDOJ 2025), and on Calder-Wang and Kim's (2024) empirical study of RealPage.²⁸

First, RealPage has a near-monopoly in providing algorithmic pricing for multifamily apartment buildings. Following RealPage's acquisition of a competing product in 2017, it is estimated that RealPage has over 95% of the algorithmic pricing market in this industry (Calder-Wang and Kim, p. 2). In addition, while the apartment rental industry is relatively fragmented and algorithmic pricing penetration is nowhere near 100%, it is reasonable to view RealPage clients as large players with market power. Calder-Wang and Kim report that the 20 largest US apartment management companies—each of which manages over 32,000 units—all use RealPage; and the DOJ complaint reports that five or fewer landlords manage a majority of all multifamily apartment units in 445 US ZIP codes. While RealPage clients account for only 10.4% of the overall US apartment rental market, they account for 21% of multifamily rentals, over 30% of buildings with over 20 units, and 34% of total units. In some markets, RealPage clients' market share is considerably higher. Calder-Wang and Kim's Table 3 reports that in 18% of market segments (consisting of a geographical submarket and a building class), a majority of multifamily apartment buildings use algorithmic pricing. In addition, their Table 5 reports that penetration exceeds 66% in many submarkets (e.g., 68% in Arapahoe County, Denver; 74% in Far Northwest, Austin; 92% in Irvine, Orange County). The DOJ complaint also contains anecdotal evidence that RealPage clients perceive themselves to have market power and the ability to impact competitors' behavior. For example, one landlord wrote, "Our very first goal we came out with immediately out of the gate is that we will not be the reason any particular sub-market takes a rate dive. So for us the strategy was to hold steady and to keep an eye on the communities around us and our competitors," (pp. 4-5). Finally, both the DOJ complaint and Calder-Wang and Kim contain extensive evidence that RealPage's price recommendations are oriented more toward joint profit-maximization than individual profit-maximization. For example,

²⁸These are the main publicly available information sources on RealPage. In turn, to our knowledge, the RealPage case is the algorithmic collusion facilitation case where the most information is currently available. Information on such cases is relatively scarce because algorithmic collusion facilitation is a new issue and because several cases are currently in litigation.

the complaint reports that “RealPage frequently tells prospective and current clients that a ‘rising tide raises all ships.’ A RealPage revenue management vice president explained that this phrase means that ‘there is greater good in everybody succeeding versus essentially trying to compete against one another in a way that actually keeps the industry down,’” (p. 12). Consistent with this anecdotal evidence, Calder-Wang and Kim estimate that RealPage adopters’ pricing fits a model of joint profit-maximization better than individual profit-maximization.²⁹ So, with the important caveat that we do not model competition with non-adopters or firms’ decision to adopt algorithmic pricing, our model seems to fit this market reasonably well.

Second, RealPage has better demand information than its clients; it does not make this information available to them; and a plausible reason why is to make it easier to control their pricing decisions. RealPage’s business is based on its superior, non-public data: e.g., it claims that it “does not have any true competitors, mainly because our data is based on real lease transaction data,” and that “we have [the] most data and the best data” (USDOJ 2024a, p. 5, 14). It obtains this data by collecting and combining data from its many clients (the conduct at the heart of its business model and the many antitrust complaints it has attracted), but also by direct private market research. As the DOJ reports (p. 11), “RealPage has an additional, complementary product called Market Analytics. Market Analytics compiles data from over 50,000 monthly phone calls that RealPage makes to landlords across the country. On these calls RealPage collects nonpublic, competitively sensitive information by floor plan on occupancy rates, effective rents, and concessions... These market surveys cover over 11 million units and approximately 52,000 properties. Landlords, including but not limited to those that use AIRM, YieldStar, or other RealPage products knowingly share this nonpublic information with RealPage.” One indication of the potential importance of this data is that DOJ’s proposed settlement bans RealPage from using competitors’ prices in the last 12 months as well as these non-public market surveys to determine recommended prices.

RealPage uses its data to recommend prices to its clients (indeed, this is the core busi-

²⁹We take no position on whether the evidence that RealPage maximizes joint rather than individual profit is fully compelling from an economic standpoint, let alone a legal one. We only claim that conceptualizing RealPage as maximizing joint profit is well-motivated.

ness that drew regulatory scrutiny), but it does not make the data itself available to its clients. Furthermore, RealPage does not provide more detailed data even if clients explicitly request it. The stated reason for this practice appears to be to reduce antitrust risk, rather than preventing price deviations as in our model. As the DOJ reports (p. 48), “When a property owner asked for information on specific competitors, Landlord 3’s director of revenue management replied that the requested tool, RealPage’s Performance Analytics with Benchmarking, did not provide information on specific competitors. The reason? Performance Analytics with Benchmarking ‘tracks transactional information therefore due [to] the potential pricing collusion, it’s anonymize[d] by RealPage.’”

However, RealPage takes a number of steps to prevent price deviations by its clients, and it is very plausible that withholding detailed demand information is helpful in this regard. RealPage’s efforts to control its clients’ pricing are well-documented and are a centerpiece of the DOJ complaint and subsequent settlement. As the complaint puts it, “[RealPage]’s recommendations are more than just ‘recommendations,’” (p. 3), because “RealPage uses multiple mechanisms to increase compliance with price recommendations,” (p. 23). These include several mechanisms that make accepting recommendations easier than declining them, including encouraging the use of an “Auto Accept” feature; allowing managers to accept recommendations in bulk but requiring them to decline recommendations one by one; and elevating attempted price overrides to a RealPage “pricing advisor,” who may further escalate the disagreement to a regional manager. According to one RealPage client, RealPage’s design is “trying to persuade [clients] to take the recommendations (almost like we made it hard to do anything but),” (p. 24).³⁰ Moreover, it is unsurprising that RealPage has to take such steps to control its clients’ pricing, because there is also evidence that RealPage’s recommended prices are not individually optimal for its clients. An example here is RealPage’s price “guardrails,” which asymmetrically favors high prices: as the complaint puts it (p. 52), “AIRM and YieldStar recommend price increases generated by the model. But the guardrails reduce or eliminate certain proposed price decreases even though the model had determined such deviations may contravene the landlord’s individual economic

³⁰It would be straightforward to incorporate these mechanisms in our model by assuming that a firm must pay an “override cost” to set a price other than the recommended one. This would just relax firm’s incentive constraints by an amount equal to the override cost, leaving our qualitative results unchanged.

interest.”³¹ In addition, Calder-Wang and Kim’s finding that RealPage adopters’ pricing fits a model of joint profit-maximization better than individual profit-maximization implies that RealPage clients have individual incentives to override RealPage’s recommendations.

Withholding information from its clients is another means through which RealPage can prevent them from overriding its recommendations. The logic of the revelation principle (Myerson, 1982) implies that communicating any information to its clients beyond their recommended prices can only tempt the clients to override RealPage’s recommendations. We lack direct evidence that RealPage consciously withholds information from its clients *in order to* reduce their temptation to override. However, such evidence is available in other settings where a cartel facilitator manages firms’ information to support collusion.³² Overall, it is clear that RealPage takes multiple measures to prevent its clients from overriding its recommendations; that RealPage has detailed demand information that it withholds from its clients; and that withholding this information reduces its clients’ potential gains from overriding its recommendations and thus facilitates their acceptance.

Finally, we can ask whether the resulting pricing patterns over the business cycle fit the predictions of our model. We focus on our central prediction that algorithmic intermediaries disclose more precise information in economic downturns—leading to lower and more flexible prices—and coarser information in booms—leading to higher and more rigid prices. The evidence in the DOJ complaint and Calder-Wang and Kim on the role of algorithmic pricing over the business cycle is broadly consistent with this pattern. First, one of Calder-Wang and Kim’s main findings is that the calendar-year treatment effect of RealPage adoption on price is strongly negative during the financial crisis and modestly positive in other years (see their Figure 7). This is consistent with viewing RealPage as providing information in accordance with our theory, where, relative to non-adopters, adopters receive precise bad news about demand in down markets (leading to sharply lower prices) and receive coarse good

³¹The above features that encourage RealPage clients to accept its recommendations and that asymmetrically favor price increases were all banned as part of the proposed DOJ settlement.

³²For example, Harrington (2006) reports that at meetings of the European isostatic graphic cartel, the firms’ representatives passed around a calculator where each representative secretly entered their own sales volume, so that ultimately only the sum of the reported sales was observable, which let each firm calculate their own market share but not their competitors’. For more on this and other related examples, see Sugaya and Wolitzky (2018).

news about demand in normal and strong markets (leading to modestly higher prices).³³ In addition, Calder-Wang and Kim’s structural estimation suggests that, relative to individual optimization, non-adopters price too high in downturns and too low in booms, which is consistent with non-adopters being less informed about demand than adopters.³⁴

In addition, the DOJ complaint provides evidence that RealPage’s data and analytics play an especially important role in price-setting during downturns. The complaint asserts that, “in down markets... [RealPage] instills pricing discipline in landlords, curbing normal fully independent competitive reactions by substituting them with interdependent decision-making,” (p. 47). Indeed, RealPage itself advertised that the “AI and the robust data in the RealPage ecosystem” helps its clients “avoid the race to the bottom in down markets,” (p. 46) and “curbs [clients’] instincts to respond to down-market conditions by either dramatically lowering price or by holding price when they are losing velocity and/or occupancy,” (p. 47)—moreover, given that Calder-Wang and Kim find a strong negative treatment effect of RealPage adoption on price in down markets, these quotations should be interpreted as emphasizing RealPage clients’ high prices in down markets relative to competitive pricing, not relative to collusion without RealPage’s demand information. RealPage also emphasizes the particular value of its price recommendations in down markets, arguing that, without its recommendations, “you’ll be pricing your renewals in the dark without insight into actual lease transaction data that [YieldStar] uses to help you make pricing decision. This is critical to price renewals right[,] especially in a downturn,” (p. 14). These quotations are all consistent with RealPage providing more precise information in down markets. In sum, while the available evidence is admittedly limited, both the quantitative evidence in Calder-Wang and Kim and the qualitative evidence in the DOJ complaint are broadly consistent with our prediction that RealPage provides more precise demand information and recommends more flexible prices in downturns than in booms.

While this section has focused on RealPage, our theory can also potentially speak to a

³³This comparison assumes that both RealPage adopters and non-adopters sustain markups relative to the static Nash benchmark, differing primarily in how markups respond to demand conditions. This viewpoint is suggested by Calder-Wang and Kim’s reduced-form results, which find strongly time-varying RealPage adoption effects but little to no time-averaged adoption effect.

³⁴Calder-Wang and Kim capture this effect with a time-varying “sophistication” parameter, which can alternatively be interpreted as reflecting firms’ information.

number of other cases. Cases concerning algorithmic pricing of hotel rooms (e.g., IDeaS and Rainmaker) raise very similar issues to RealPage. Algorithmic pricing of retail gasoline also raises similar issues, although the empirical literature on this topic to date (e.g., Assad et al., 2024) has not emphasized business cycle fluctuations. There are also some recent cases concerning algorithmic cartel facilitation where the central issue appears to be facilitating monitoring of competitors' actions rather than demand conditions, including cases involving the data analytics companies AgriStats and Circana, which provide detailed production and price information for the poultry and pork markets (AgriStats) and the frozen potato market (Circana). Finally, another setting that raises similar issues is information-sharing in bidding rings in procurement auctions. A specific example comes from the bidding ring organized by the Kumatori Contractors Cooperative, studied by Kawai, Nakabayashi, and Ortner (2024). This ring took drastic steps to limit bidders' information about the cost of completing (only) the largest construction project they bid on: “[The director of the Cooperative] told the members that he would be collecting, from each of the invited bidders, the detailed project plan that the town distributes at the on-site briefing. This was understood by the members of the Cooperative as a preventative measure to make defection more difficult by making it harder for other firms to estimate costs,” (ibid.). This matches our result that cartels censor information in those states of the world where deviating is most tempting.³⁵

6 Extensions

We consider two extensions of our baseline model. Section 6.1 lets the state follow a Markov process rather than being iid across time. Section 6.2 assumes a linear demand system rather than homogeneous goods.

³⁵A subtlety in mapping this example to our model is that the “tempting state” is a large procurement project, which has both a higher (observed) reserve price and a higher (unobserved) completion cost than a typical project. To capture this example, one can extend the model to allow both an observable stochastic component (the reserve price) and an unobserved one (the cost).

6.1 Persistent Demand or Cost

We now consider the case where the state follows a Markov process: we assume that the current state s' is drawn from a distribution F_s , where s is the previous period's state. This realistic extension of the baseline iid model illustrates how our results generalize and also yields some new insights. The analysis of this section is inspired by Haltiwanger and Harrington (1991), Kandori (1992), and Bagwell and Staiger (1997), who extended Rotemberg and Saloner's (1986) iid model to various Markov processes.

To accommodate the Markov case, we need to preserve the property that expected current and future profit is measurable with respect to current mean demand (or cost, but we continue to focus on the demand case). This requires two assumptions. First, we assume that the current demand state is revealed at the end of each period, so the algorithm does not carry private information across periods. This is realistic if firms observe their sales at the end of each period. Second, we assume that the Markov transition rule F_s is affine in s , so the distribution over tomorrow's state depends only on today's mean state:

$$F_s(s') = \frac{\bar{s} - s}{\bar{s} - \underline{s}} F_{\underline{s}}(s') + \frac{s - \underline{s}}{\bar{s} - \underline{s}} F_{\bar{s}}(s') \quad \text{for all } s, s'.$$

For example, F_s is affine in s when there is a binary underlying demand state \mathbf{s} and s is a continuous signal of \mathbf{s} satisfying $\Pr(\mathbf{s} = \bar{s}|s) = (s - \underline{s}) / (\bar{s} - \underline{s})$. We also assume that the distribution of s in period 1 equals F_{s_0} for some $s_0 \in [\underline{s}, \bar{s}]$.

Affineness allows both *positive persistence*—where $F_{\bar{s}}$ first-order stochastically dominates $F_{\underline{s}}$ —and *negative persistence*—where $F_{\underline{s}}$ first-order stochastically dominates $F_{\bar{s}}$. Both cases are of interest: positive persistence is arguably more natural, while negative persistence has been used to capture cyclical demand movements (Haltiwanger and Harrington, 1991).

The characterization of the optimal disclosure policy and collusive prices are the same as in the iid case, except that now the expected value of collusive profit $V(s)$ depends on the previous period's state s . The optimal collusive profit for each last-period demand state s must now be calculated simultaneously as the component-wise greatest fixed point

$(V^*(s))_{s \in [\underline{s}, \bar{s}]}$ of the following system of equations in s :

$$V(s) = (1 - \delta) \max_{G \in MPC(F_s)} \mathbb{E}^{x \sim G} [\min \{\Pi^m(x), \Pi^{\max}(\delta, n, \mathbb{E}^{F_x}[V(s')])\}] + \delta \mathbb{E}^{F_s}[V(s')]. \quad (5)$$

The right-hand side of (5) is bounded and increasing in $V(s')$ for all s, s' , so the greatest fixed point is well-defined by Tarski's theorem. We also define $W^*(s) = \mathbb{E}^{F_s}[V^*(s')]$, so we have

$$V^*(s) = (1 - \delta) \max_{G \in MPC(F_s)} \mathbb{E}^G [\min \{\Pi^m(x), \Pi^{\max}(\delta, n, W^*(x))\}] + \delta W^*(s) \quad \text{for all } s. \quad (6)$$

Note that, since F_s is affine in s , so is $W^*(s)$.

With persistent demand, the appropriate notion of a (symmetric) stationary strategy is that the disclosure policy G depends only the previous period's demand state, while the on-path price $p(s)$ at realized mean demand state s remains independent of the history of past demand realizations (and, in particular, is independent of the current-period disclosure policy). With this definition, Lemma 1 generalizes as follows.

Lemma 3 *The expected present value of optimal collusive profit in each state s equals $V^*(s)$ and is attained by a symmetric, stationary, grim trigger equilibrium. Moreover, a collection of disclosure policies $(G_s)_{s \in [\underline{s}, \bar{s}]}$, one for each last-period demand state s , is optimal if and only if, for each s , G_s solves the maximization problem in (5) with $V(\cdot) = V^*(\cdot)$.*

Lemma 3 reduces the problem of finding an optimal equilibrium to the family of static information design problems on the right-hand side of (5), where the function $V^*(\cdot)$ satisfies the fixed point condition.³⁶

As in the iid case, collusion is impossible if $\delta < (n - 1)/n$. Conversely, if $\delta \geq (n - 1)/n$ then monopoly profit under no disclosure given the least-favorable previous period demand state (e.g., \underline{s} in the positively persistent case; \bar{s} in the negatively persistent case), $\Pi^m(\min \{\mathbb{E}^{F_{\underline{s}}}[s], \mathbb{E}^{F_{\bar{s}}}[s]\})$, is attainable for any initial state.

³⁶The proof is a straightforward generalization of the proof of Lemma 1: the only difference is that the present value of equilibrium profits, the probability distribution over next-period demand states, and the values \bar{V} and v_i defined in the second part of the proof are all now functions of the current expected state s .

Lemma 4 *If $\delta < (n-1)/n$ then $V^*(s) = 0$ for all s . Conversely, if $\delta \geq (n-1)/n$ then $V^*(s) \geq \Pi^m(\min\{\mathbb{E}^{F_{\underline{s}}}[s], \mathbb{E}^{F_{\bar{s}}}[s]\})$ for all s .*

Proof. If $\delta < (n-1)/n$ then $\Pi^{\max}(\delta, n, V) < V$ for all $V > 0$. Let $s_0 = \operatorname{argmax}_s V^*(s)$, which is well-defined because $\Pi^m(s)$ is continuous and $W^*(s)$ is affine. Suppose for contradiction that $V^*(s_0) > 0$. Then, since $W^*(s) \leq V^*(s_0)$ for all s (as $W^*(s) = \mathbb{E}^{F_s}[V^*(s')]$), the right-hand side of (6) at $s = s_0$ is strictly less than $V^*(s_0)$, a contradiction. Hence, $V^*(s_0) = 0$, and therefore $V^*(s) = 0$ for all s .

Conversely, if $\delta \geq (n-1)/n$ then under no information disclosure it is an equilibrium to set on-path price $\min\{p : \Pi(p, \mathbb{E}^{F_s}[\tilde{s}]) = \Pi^m(\min\{\mathbb{E}^{F_{\underline{s}}}[\tilde{s}], \mathbb{E}^{F_{\bar{s}}}[\tilde{s}]\})\}$ when the previous period demand state is s (noting that this price is well-defined by the intermediate value theorem, as $\Pi(p, \mathbb{E}^{F_s}[\tilde{s}])$ is continuous in p , $\Pi(0, \mathbb{E}^{F_s}[\tilde{s}]) = 0$, and $\Pi(p^m(\mathbb{E}^{F_s}[\tilde{s}]), \mathbb{E}^{F_s}[\tilde{s}]) = \Pi^m(\mathbb{E}^{F_s}[\tilde{s}]) \geq \Pi^m(\min\{\mathbb{E}^{F_{\underline{s}}}[\tilde{s}], \mathbb{E}^{F_{\bar{s}}}[\tilde{s}]\})$) and off-path price zero. ■

We now characterize the optimal disclosure policy as a function of the last-period state s in the non-trivial case where $\delta \geq (n-1)/n$. First, let s^* solve

$$\Pi^m(s^*) = \Pi^{\max}(\delta, n, W^*(s^*)) \quad (7)$$

if such a demand state exists, and let $s^* = \bar{s}$ otherwise.³⁷ Next, for each last-period state s , let $\hat{s}(s)$ satisfy

$$\mathbb{E}^{F_s}[\tilde{s} | \tilde{s} \geq \hat{s}(s)] = s^* \quad (8)$$

if such a state exists, and let $\hat{s}(s) = \underline{s}$ otherwise. Note that, by Lemma 4 and our assumption that $\delta \geq (n-1)/n$, we have $V^*(\underline{s}) \geq \Pi^m(\min\{\mathbb{E}^{F_{\underline{s}}}[s], \mathbb{E}^{F_{\bar{s}}}[s]\})$, and hence $s^* \geq \min\{\mathbb{E}^{F_{\underline{s}}}[s], \mathbb{E}^{F_{\bar{s}}}[s]\}$, so (8) admits a solution $\hat{s}(s) \in [\underline{s}, \bar{s}]$ for $s = \operatorname{argmin}_{s \in \{\underline{s}, \bar{s}\}} \mathbb{E}^{F_s}[s']$. However, in contrast to the iid case, (8) does not always admit a solution $\hat{s}(s)$ for all last-period demand states s : in this case, the distribution F_s is so high that $\mathbb{E}^{F_s}[s'] > s^*$, in which case no disclosure of the current demand state is optimal, and the optimal price is $\min\{p : \Pi(p, \mathbb{E}^{F_s}[s']) = \Pi^{\max}(\delta, n, W^*(\mathbb{E}^{F_s}[s']))\}$, which is less than the corresponding

³⁷There is at most one solution to (7). If $W^*(s)$ is decreasing, this is immediate, as the left-hand side of (7) is increasing and the right-hand side is decreasing. If $W^*(s)$ is increasing, this follows because, since $\delta \geq (n-1)/n$, we have $\Pi^m(\underline{s}) \leq \Pi^{\max}(\delta, n, \Pi^m(\underline{s})) < \Pi^{\max}(\delta, n, \Pi^m(\min\{\mathbb{E}^{F_{\underline{s}}}[s], \mathbb{E}^{F_{\bar{s}}}[s]\})) \leq \Pi^{\max}(\delta, n, W^*(\underline{s}))$, and the left-hand side of (7) is convex while the right-hand side is linear.

monopoly price $p^m(\mathbb{E}^{F_s}[s'])$. Thus, in the Markov case, following last-period states that make firms sufficiently optimistic about the current state, the optimal collusive policy can entail no information disclosure and a price below the corresponding monopoly price.

We can now characterize the optimal disclosure and pricing policy in the Markov case.

Theorem 2 *With stochastic demand, the unique optimal disclosure policy as a function of the last-period demand state s is the upper censorship policy that discloses demand states below $\hat{s}(s)$ and conceals demand states above $\hat{s}(s)$. The optimal collusive price $p(\tilde{s}; s)$ (which is unique except when $\hat{s}(s) = \underline{s}$) when the current realized mean demand state is \tilde{s} and the last-period demand state is s is given by*

$$p(\tilde{s}; s) = \begin{cases} p^m(\tilde{s}) & \text{if } \tilde{s} < \hat{s}(s), \\ p^m(s^*) & \text{if } \tilde{s} \geq \hat{s}(s) > \underline{s}, \\ \min \{p : \Pi(p, \mathbb{E}^{F_s}[s']) = \Pi^{\max}(\delta, n, W^*(\mathbb{E}^{F_s}[s']))\} & \text{if } \hat{s}(s) = \underline{s}. \end{cases} \quad (9)$$

Moreover, under positive persistence, $\hat{s}(s)$ is decreasing, so the optimal policy discloses less information when last-period demand is higher; conversely, under negative persistence, $\hat{s}(s)$ is increasing, so the optimal policy discloses more information when last-period demand is higher.

With stochastic costs, the unique optimal disclosure policy is the analogous lower censorship policy that discloses cost states above $\hat{s}(s)$ satisfying $\mathbb{E}^{F_s}[\tilde{s} | \tilde{s} \leq \hat{s}(s)] = s^*$ and conceals cost states below $\hat{s}(s)$. The unique optimal collusive price $p(s)$ in state s is given by (9) with the reversed inequalities and \bar{s} in place of \underline{s} . Moreover, under positive persistence, $\hat{s}(s)$ is decreasing, so the optimal policy discloses more information when last-period cost is higher; conversely, under negative persistence, $\hat{s}(s)$ is increasing, so the optimal policy discloses less information when last-period cost is higher.

Proof. The proof is a straightforward generalization of the proof of Theorem 1. The main difference is that, since $W^*(s)$ is affine, the function $\min \{\Pi^m(s), \Pi^{\max}(\delta, n, W^*(s))\}$ is now “convex-then-linear” in s , rather than “convex-then-constant” as in the iid case. The same argument as in the proof of Theorem 1 implies that, when $\hat{s}(s) > \underline{s}$, upper censorship is optimal, with mean demand among concealed states equal to the point s^* where $\Pi^m(s^*) =$

$\Pi^{\max}(\delta, n, W^*(s^*))$. A similar argument shows that, when $\hat{s}(s) = \underline{s}$, no disclosure is optimal, with a price p being the smallest price to satisfy $\Pi(p, \mathbb{E}^{F_s}[s']) = \Pi^{\max}(\delta, n, W^*(\mathbb{E}^{F_s}[s']))$. Finally, it is immediate from (8) that $\hat{s}(s)$ is decreasing under positive persistence and increasing under negative persistence. ■

The new insights of Theorem 2 concern how optimal disclosure depends on last-period demand. When the optimal disclosure policy is non-trivial (i.e., $\hat{s}(s) \in (\underline{s}, \bar{s})$), so some states are disclosed and others are censored), the mean censored state is fixed at s^* , regardless of the last-period demand state s . With positive persistence, this requires greater censoring (lower \hat{s}) when the last-period demand state is higher: intuitively, the algorithm discloses less information following good periods, when firms are optimistic about current demand and are thus more tempted to deviate. Conversely, with negative persistence, the algorithm discloses more information following good periods, when firms are pessimistic and are thus less tempted to deviate.

In contrast to the iid case, for certain last-period states it can be optimal for the algorithm to disclose no information and recommend prices below the corresponding monopoly price.³⁸ For example, with positive persistence, it can be optimal to fully reveal current demand when last-period demand was low (so $\hat{s}(s) = \bar{s}$), partially reveal current demand when last-period demand was intermediate (so $\hat{s}(s) \in (\underline{s}, \bar{s})$), and reveal nothing about current demand when last-period demand was high (so $\hat{s}(s) = \underline{s}$)—moreover, in the last case, the optimal price $p(\mathbb{E}^{F_s}[s'] ; s)$ satisfies $\Pi(p, \mathbb{E}^{F_s}[s']) = \Pi^{\max}(\delta, n, W^*(\mathbb{E}^{F_s}[s']))$, and so is less than the monopoly price $p^m(\mathbb{E}^{F_s}[s'])$ and can even be decreasing in s . Thus, while collusive prices are always monotone in current demand (as in the iid case and in contrast to Rotemberg Saloner), they may be non-monotone in last-period demand. Notably, the expected price conditional on last-period demand can be single-peaked, a result that recovers some of Rotemberg and Saloner’s intuition.³⁹

In addition to these novel points, Theorem 2 shows that the main results from the iid case generalize to the Markov case. In particular, for any last-period demand state s , optimal collusion entails price rigidity at high current demand states, supra-monopoly prices over an

³⁸Recall that in the iid case, no disclosure is only optimal in the knife-edge case where $\delta = (n - 1)/n$.

³⁹Whether prices actually display this pattern depends on whether $\Pi(p, \tilde{s})$ or $\Pi^{\max}(\delta, n, W^*(\tilde{s}))$ increases faster in $\tilde{s} = \mathbb{E}^{F_s}[s']$ over the range $\{s : \hat{s}(s) = \underline{s}\}$.

intermediate range of states, and more flexible prices when n is lower, δ is higher, or F_s is more informative.

Finally, in Appendix D, we provide a numerical example showing that the effect of greater demand persistence on collusive profit, consumer surplus, and the amount of information disclosure can all be non-monotone.

6.2 Differentiated Products

We now extend the baseline iid model from the homogeneous-goods setting of Rotemberg and Saloner to a symmetric, linear demand system. This is a workhorse demand system for studying oligopoly pricing under uncertainty (e.g., Vives, 2001, Chapter 8) and has recently been used by Harrington (2022, 2025) to study oligopoly pricing with a third-party pricing algorithm. In Appendix A, we further extend the analysis of this section to more general demand systems, which include as a special case the baseline, homogeneous goods model, with or without capacity constraints.

In this section, we assume that there exists a constant $b \in [0, 1/(n-1)]$ such that firm i 's payoff at price vector $\mathbf{p} = (p_1, \dots, p_n)$ and state s equals

$$\pi_i(\mathbf{p}, s) = p_i \left(s + b \sum_{j \neq i} p_j - p_i \right)$$

for the stochastic demand case, or

$$\pi_i(\mathbf{p}, s) = (p_i - s) \left(1 + b \sum_{j \neq i} p_j - p_i \right)$$

for the stochastic cost case.⁴⁰ Note that $\pi_i(\mathbf{p}, s)$ remains affine in s . For concreteness, we focus on the stochastic demand case.

To analyze this model, we first define the monopoly price and the per-firm monopoly

⁴⁰The condition $b \geq 0$ implies that the firms' products are substitutes. The condition $b \leq 1/(n-1)$ implies that profits are bounded and is satisfied whenever the demand system results from utility maximization by a representative consumer (Amir, Erickson, and Jin, 2017).

profit at a public mean belief s as

$$p^m(s) = \frac{s}{2(1 - (n-1)b)} \quad \text{and} \quad \pi^m(s) = \frac{s^2}{4(1 - (n-1)b)},$$

and we define the static Nash equilibrium price and the per-firm static Nash profit at a public mean belief s as

$$p^N(s) = \frac{s}{2 - (n-1)b} \quad \text{and} \quad \pi^N(s) = \left(\frac{s}{2 - (n-1)b} \right)^2.$$

Notice that $\pi^N(s)$ is convex in s , which implies that the expected static Nash profit $\mathbb{E}^G[\underline{\pi}(s)]$ is minimized over $G \in MPC(F)$ by taking $G = \delta_{\mathbb{E}^F[s]}$: that is, by entirely concealing the state. Thus, the *no-information Nash outcome*, where the algorithm discloses no information and expect profit equals

$$\underline{\pi} = \pi^N(\mathbb{E}^F[s]),$$

minimizes the firms' static Nash profit.

We characterize the pure-strategy, subgame perfect equilibrium that maximizes collusive profit among all equilibria where off-path expected profit equals $\underline{\pi}$. Note that, in contrast to the homogeneous goods case where the static Nash profit equals the minimax payoff of 0, the off-path payoff $\underline{\pi}$ is strictly greater than the minimax payoff of $\mathbb{E}^F[s]^2/4$. We are thus focusing on equilibria sustained by the threat of *Nash reversion*. As is well-known, this entails a loss of optimality in the class of all pure-strategy, subgame perfect equilibria. To find the optimal equilibrium in this larger class, one would simultaneously find the worst equilibrium for each firm as a fixed point, following Abreu (1988). However, except for giving a different value for the off-path payoff $\underline{\pi}$, this procedure would yield the same characterization of optimal equilibrium prices and information disclosure on path. Our qualitative results are thus insensitive to the specification of off-path payoffs.

To find the optimal equilibrium (in the Nash reversion class), denote per-firm profit when all firms set price p in demand state s by

$$\pi(p, s) = p(s + (n-1)bp) - p^2,$$

and denote a firm's maximum payoff from a deviation when all firms set price p in demand state s by

$$\pi^d(p, s) = \left(\frac{s + (n-1)bp}{2} \right)^2.$$

Next, for any $v \geq 0$, let $p^{\max}(s, v)$ be the larger solution to the quadratic equation

$$\pi^d(p, s) - \pi(p, s) = \frac{\delta}{1-\delta}(v - \underline{\pi}),$$

which is given by

$$p^{\max}(s, v) = \frac{s + 2\sqrt{\frac{\delta}{1-\delta}(v - \underline{\pi})}}{2 - (n-1)b},$$

and let

$$\pi^{\max}(s, v) = \pi(s, p^{\max}(s, v)) = \frac{s^2 + 2(n-1)bs\sqrt{\frac{\delta}{1-\delta}(v - \underline{\pi})} - 4(1 - (n-1)b)\frac{\delta}{1-\delta}(v - \underline{\pi})}{(2 - (n-1)b)^2}.$$

Thus, $p^{\max}(s, v)$ is the greatest incentive compatible price in state s in a symmetric equilibrium with on-path per-firm continuation payoff v and off-path per-firm profit $\underline{\pi}$, and $\pi^{\max}(s, v)$ is the corresponding per-firm profit in state s . While the details of these formulas will not matter, it is important to note that $p^{\max}(s, v)$ is increasing in s and $\pi^{\max}(s, v)$ is increasing and convex in s (in addition to depending implicitly on the parameters δ , n , and b). This contrasts with the homogeneous goods case, where $p^{\max}(s, v)$ is decreasing in s and $\pi^{\max}(s, v)$ is constant in s for $s > s^*$.⁴¹

Now, following the homogeneous goods case, define v^* as the greatest fixed point of the equation

$$v = \max_{G \in MPC(F)} \mathbb{E}^G [\min \{\pi^m(x), \pi^{\max}(x, v)\}]. \quad (10)$$

(Here, we use lower-case letters for per-firm payoffs, whereas in the baseline model we used capital letters for the corresponding industry payoffs.) Lemma 1 extends to the current

⁴¹Intuitively, with undifferentiated goods, $\pi^d(p, s) = n\pi(p, s)$, so $\pi^d(p, s) - \pi(p, s)$ is a constant multiple of $\pi(p, s)$, and hence an upper bound for $\pi^d(p, s) - \pi(p, s)$ (independent of s) implies an upper bound for $\pi(p, s)$ (independent of s). In contrast, with differentiated goods, $\pi^d(p, s) - \pi(p, s)$ is not a constant multiple of $\pi(p, s)$, so an upper bound for $\pi^d(p, s) - \pi(p, s)$ (independent of s) implies an upper bound for $\pi(p, s)$ that depends on s .

setting: in particular, optimal collusive profit per-firm equals v^* . Formally, this follows from the more general Lemma 5 in Appendix A. The key step in the proof is that symmetric pricing remains optimal in each demand state. This holds because, for any price vector $\mathbf{p} = (p_1, \dots, p_n)$, we have

$$\pi \left(\frac{1}{n} \sum_i p_i, s \right) \geq \frac{1}{n} \sum_i \pi_i(\mathbf{p}, s) \quad \text{and} \quad \pi^d \left(\frac{1}{n} \sum_i p_i, s \right) \leq \frac{1}{n} \sum_i \pi_i^d(\mathbf{p}_{-i}, s)$$

(by straightforward calculation), so replacing any asymmetric price vector \mathbf{p} with the symmetric vector $(\sum_i p_i/n, \dots, \sum_i p_i/n)$ increases profits without violating incentive constraints.

It remains to solve the static information design problem, (10). Let s^* satisfy $p^m(s^*) = p^{\max}(s^*, v^*)$ (or, equivalently, $\pi^m(s^*) = \pi^{\max}(s^*, v^*)$), so that

$$s^* = \frac{4(1 - (n-1)b) \sqrt{\frac{\delta}{1-\delta} (v^* - \underline{\pi})}}{(n-1)b}.$$

Note that $d\pi^m(s)/ds$ is greater than the right-derivative of $\pi^{\max}(s, v^*)$ with respect to s at $s = s^*$, because $\pi^{\max}(s, v^*)$ incorporates an additional constraint at $s = s^*$. Therefore, the kink in the objective function $\min \{\pi^m(s), \pi^{\max}(s, v)\}$ at $s = s^*$ is concave, so the function $\min \{\pi^m(s), \pi^{\max}(s, v)\}$ is piecewise-convex but not globally convex.

We now define a pair of states (\hat{s}_L, \hat{s}_H) as follows. First, if $s^* < \underline{s}$, define $\hat{s}_L = \hat{s}_H = \underline{s}$. Second, if $s^* \in [\underline{s}, \mathbb{E}^F[s]]$, define (\hat{s}_L, \hat{s}_H) so that

$$\begin{aligned} \hat{s}_L &< \hat{s}_H, \\ \mathbb{E}^F[s | s \in [\hat{s}_L, \hat{s}_H]] &= s^*, \quad \text{and} \\ \frac{\hat{s}_H - s^*}{\hat{s}_H - \hat{s}_L} \pi^m(\hat{s}_L) + \frac{s^* - \hat{s}_L}{\hat{s}_H - \hat{s}_L} \pi^{\max}(\hat{s}_H, v^*) &= \pi^m(s^*), \end{aligned} \tag{11}$$

if such a pair exists, and otherwise define $\hat{s}_L = \underline{s}$ and define \hat{s}_H so that $s^* = \mathbb{E}^F[s | s \leq \hat{s}_H]$. Third, if $s^* \in [\mathbb{E}^F[s], \bar{s}]$, define (\hat{s}_L, \hat{s}_H) so that (11) holds if such a pair exists, and otherwise define \hat{s}_L so that $s^* = \mathbb{E}^F[s | s \geq \hat{s}_L]$ and define $\hat{s}_H = \bar{s}$. Finally, if $s^* > \bar{s}$, define $\hat{s}_L = \hat{s}_H = \bar{s}$.

Note that (\hat{s}_L, \hat{s}_H) is well-defined, because at most one pair (\hat{s}_L, \hat{s}_H) satisfies (11). This can be seen geometrically in Figure 2: since $\pi^m(s)$ and $\pi^{\max}(s, v^*)$ are convex in s , if (\hat{s}_L, \hat{s}_H)

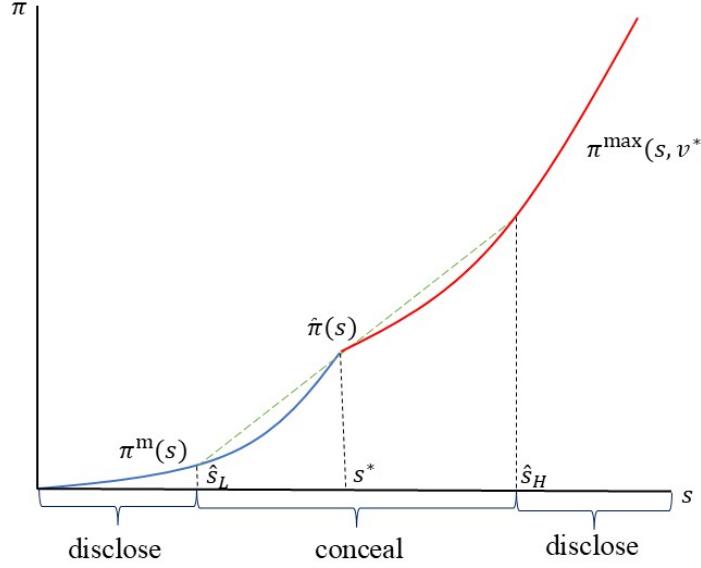


Figure 2: The optimal disclosure policy with a symmetric, linear demand system. The cutoffs (\hat{s}_L, \hat{s}_H) are defined in the text. The optimal information policy discloses demand states $s < \hat{s}_L$ and recommends the corresponding monopoly price, $p^m(s)$; conceals demand states $s \in [\hat{s}_L, \hat{s}_H]$ and recommends the monopoly price conditional on this information, $p^m(s^*)$; and reveals demand states $s > \hat{s}_H$ and recommends the maximum incentive compatible price, $p^{\max}(s, v^*)$. The auxiliary objective function $\tilde{\pi}(s)$ is defined in the proof of Theorem 4 in Appendix A.

satisfies (11) and $\mathbb{E}^F[s | s \in [s_L, s_H]] = s^*$ for $s_L > \hat{s}_L$ and $s_H < \hat{s}_H$, then $\frac{s_H - s^*}{s_H - s_L} \pi^m(s_L) + \frac{s^* - s_L}{s_H - s_L} \pi^{\max}(s_H, v^*) < \pi^m(s^*)$; while if $\mathbb{E}^F[s | s \in [s_L, s_H]] = s^*$ for $s_L < \hat{s}_L$ and $s_H > \hat{s}_H$, then $\frac{s_H - s^*}{s_H - s_L} \pi^m(s_L) + \frac{s^* - s_L}{s_H - s_L} \pi^{\max}(s_H, v^*) > \pi^m(s^*)$.

The following is the main result of this section. It is an implication of the more general Theorem 4 in Appendix A.

Theorem 3 *With stochastic demand and differentiated products, the unique optimal disclosure policy discloses demand states below \hat{s}_L and above \hat{s}_H and conceals demand states in the interval $[\hat{s}_L, \hat{s}_H]$. The unique optimal collusive price $p(s)$ in state s is given by*

$$p(s) = \begin{cases} p^m(s) & \text{if } s < \hat{s}_L, \\ p^m(s^*) & \text{if } s \in [\hat{s}_L, \hat{s}_H], \\ p^{\max}(s, v^*) & \text{if } s > \hat{s}_H. \end{cases}$$

Theorem 3 is illustrated in Figure 2. Mathematically, since the objective function $\min \{\pi^m(s), \pi^{\max}(s, v^*)\}$ is piecewise-convex with a concave kink at s^* , it is optimal to

disclose the lowest and highest demand states while concealing intermediate states. (In general, any of these three intervals can be empty, so that full disclosure, no disclosure, upper censorship, or lower censorship can also be optimal.) The intuition is that with differentiated goods, incentive constraints continue to bind in high demand states (i.e., at $s > s^*$), but there is no longer an upper bound for the maximum on-path equilibrium profit (in contrast to the homogeneous goods case), and indeed maximum profit is convex in s on the region $s > s^*$. Consequently, it is optimal to disclose the highest demand states while cutting price to satisfy incentive constraints.

The optimal equilibrium with differentiated goods described in Theorem 3 shares several key features with the homogeneous goods case. In particular, it is optimal to censor a region of demand states while recommending a rigid price that is above the monopoly price for an interval of states. However, it is no longer always optimal to censor the highest demand states: depending on parameters, it could be optimal to censor an intermediate interval of states, or only the lowest states. In addition, if the highest demand states are disclosed, collusive prices in these states are below the corresponding monopoly prices.

7 Conclusion

This paper has developed a tractable model of an intermediary that possesses information on market demand or the cost of serving the market that is superior to that of the firms competing for the market and that selectively discloses this information to maximize the firms' profit in the best collusive equilibrium. Our main motivation is the rise of third-party pricing algorithm providers such as RealPage in apartment rentals, A2i Systems and Kali-brate in retail gasoline, and IDEaS and Rainmaker in hotel rooms, but the theory applies equally to any cartel facilitator that controls the participating firms' information. We adapt the canonical Rotemberg Saloner (1986) model of repeated Bertrand competition with stochastic demand by letting an intermediary selectively disclose demand or cost information. Assuming that expected profit is determined by the expected state, we show that with homogeneous goods, optimal information disclosure takes an *upper censorship* form: demand states s below a cutoff \hat{s} are disclosed and result in the corresponding monopoly price $p^m(s)$,

while demand states above \hat{s} are concealed and result in the monopoly price for the mean concealed state, $p^m(\mathbb{E}[s|s > \hat{s}])$. The resulting pricing policy entails *price rigidity*, as well as *supra-monopoly prices* for a range of intermediate demand states. Prices are more flexible when the market is more concentrated, the firms are more patient, or the algorithm is more accurate. In turn, price flexibility reduces expected consumer surplus whenever improving a monopoly's information does so. This result suggests that improved algorithmic demand prediction is likely to reduce expected consumer surplus, in contrast to prior studies that find more optimistic results when the algorithm always discloses its predictions (Sugaya and Wolitzky, 2018; Miklós-Thal and Tucker, 2019). Finally, most results survive in more general specifications with demand persistence, product differentiation, or capacity constraints, although the specific form of the optimal censorship policy depends on the demand system, and collusive prices can sometimes fall short of the corresponding monopoly prices.

Many of our assumptions can be further relaxed at the cost of a more intricate analysis. First, if expected profit depended on the entire distribution of the unknown state rather than only its mean, we would have a non-linear information design problem, where disclosure policies that pool intervals of states together (like upper censorship) are typically sub-optimal (Kolotilin, Corrao, and Wolitzky, 2024). However, upper censorship is approximately optimal if the information design problem is close to linear.⁴² Second, if the intermediary's objective differs from maximizing industry profit or if some firms do not use the intermediary, as in Harrington (2022), the model must be extended to incorporate the intermediary's incentives and firms' incentives to use the intermediary. Third, if the colluding firms are asymmetric, there is no longer a unique cartel-optimal equilibrium, and one must instead analyze the Pareto frontier of the equilibrium payoff set and consider non-stationary equilibria as in Abreu (1986). Fourth, in practice, algorithmic intermediaries can also facilitate collusion by systematizing monitoring of firms' prices. This could be incorporated by considering an imperfect monitoring version of our model. Finally, allowing asymmetric, private information disclosure by the algorithm—or exogenous asymmetric, private information for the firms, which could possibly be elicited by the algorithm—appears challenging but potentially

⁴²It is also likely that, if the problem is close to linear, every optimal disclosure policy approximates upper censorship. A result along these lines in a related class of information design problems is Theorem 3 of Kolotilin and Wolitzky (2025).

realistic and insightful.⁴³ These are all interesting directions for future research.

A General Payoffs and Capacity Constraints

In this appendix, we derive a version of our main result for a general class of payoff functions that includes both the undifferentiated and differentiated goods models analyzed in the text, as well as a model of undifferentiated goods with capacity constraints.

As in Section 6.2, let $\pi_i(\mathbf{p}, s)$ denote firm i 's profit at price vector $\mathbf{p} = (p_1, \dots, p_n)$ and demand state s .⁴⁴ We also define firm i 's maximum deviation profit at price vector \mathbf{p} by

$$\pi_i^d(\mathbf{p}, s) = \max_{p_i} \pi_i(p_1, \dots, p_{i-1}, p_i, p_{i+1}, \dots, p_n, s).$$

We impose four assumptions on payoffs. First, we assume that payoffs are symmetric, continuous, and quasi-concave in a firm's own price.⁴⁵

Assumption 1 For any price vector (p_1, \dots, p_n) , state s , firm i , and permutation ϕ on $\{1, \dots, n\}$, we have $\pi_i(p_{\phi(1)}, \dots, p_{\phi(n)}, s) = \pi_{\phi(i)}(p_1, \dots, p_n, s)$. In addition, $\pi_i(\mathbf{p}, s)$ is continuous in \mathbf{p} and quasi-concave in p_i for all \mathbf{p}_{-i} .

Second, we assume that $\pi_i(\mathbf{p}, s)$ is affinely increasing in s .

Assumption 2 For any price vector \mathbf{p} and firm i , we have $\pi_i(\mathbf{p}, \underline{s}) \leq \pi_i(\mathbf{p}, \bar{s})$ and

$$\pi_i(\mathbf{p}, s) = \frac{\bar{s} - s}{\bar{s} - \underline{s}} \pi_i(\mathbf{p}, \underline{s}) + \frac{s - \underline{s}}{\bar{s} - \underline{s}} \pi_i(\mathbf{p}, \bar{s}) \quad \text{for all } s.$$

Assumptions 1 and 2 imply that, for any mean public belief s , there exists a symmetric, pure-strategy, static Nash equilibrium \mathbf{p} . Let $\underline{\pi}$ denote the lowest payoff from such a Nash

⁴³The former question would combine the forces in the present paper with those in Ortner, Sugaya, and Wolitzky (2024). The latter would concern optimal Bayes correlated equilibria (as in, e.g., Smolin and Yamashita, 2025) or communication equilibria (as in, e.g., Goltzman and Pavlov, 2014).

⁴⁴This appendix focuses on stochastic demand. The stochastic cost case is analogous.

⁴⁵The symmetry notion in Assumption 1 is known as *total symmetry*. The results in this appendix also hold under the weaker notion of *weak symmetry*: for any pair of firms i and j , there exists a permutation ϕ on $\{1, \dots, n\}$ such that $\pi_i(i) = j$ and, for any price vector (p_1, \dots, p_n) , state s , and firm k , we have $\pi_i(p_{\phi(1)}, \dots, p_{\phi(n)}, s) = \pi_{\phi(k)}(p_1, \dots, p_n, s)$. The examples in the current paper are all totally symmetric, but some oligopoly models (e.g., Salop's circle model) are only weakly symmetric. See Plan (2023).

equilibrium at mean public belief $\mathbb{E}^F[s]$. As in Section 6.2, for concreteness we restrict attention to pure-strategy equilibria where off-path payoffs are given by $\underline{\pi}$.

Next, we define a firm's profit when all firms set price p in demand state s as $\pi(p, s) = \pi_i(p, \dots, p, s)$, and we define the corresponding maximum deviation profit as $\pi^d(p, s) = \pi_i^d(p, \dots, p, s)$. Our third assumption is that for any price vector, there exists a constant price vector that weakly increases industry profit without increasing the average of the firms' deviation gains.

Assumption 3 For any price vector \mathbf{p} and state s , there exists a price $p \geq 0$ such that

$$\pi(p, s) \geq \frac{1}{n} \sum_i \pi_i(\mathbf{p}, s) \quad \text{and} \quad \pi^d(p, s) - \pi(p, s) \leq \frac{1}{n} \sum_i (\pi_i^d(\mathbf{p}, s) - \pi_i(\mathbf{p}, s)).$$

Finally, we assume that the profit function $\pi(p, s)$ is quasi-concave in p .

Assumption 4 $\pi(p, s)$ is quasi-concave in p with a well-defined monopoly profit $\pi^m(s) = \pi(p^m(s), s) = \max_p \pi(p, s)$ for each s .

Assumptions 1–4 are all satisfied in both the baseline undifferentiated goods model and the linear differentiated goods model of Section 6.2. Another important example is undifferentiated goods with capacity constraints, where

$$\pi_i(\mathbf{p}, s) = \mathbf{1} \left\{ p_i = \min_j p_j \right\} \times p_i \min \left\{ \frac{D(\min_j p_j, s)}{|\{k : p_k = \min_j p_j\}|}, C \right\},$$

where $D(p, s)$ is the industry demand and C is a per-firm capacity constraint. This specification satisfies Assumptions 1, 3, and 4, and it also satisfies Assumption 2 if s is binary.⁴⁶

We now characterize the optimal collusive profit level, information disclosure policy, and equilibrium in this general model. First, for any $v \geq 0$, let $p^{\max}(s, v)$ be the greatest solution p in $[0, p^m(s)]$ to the equation

$$\pi^d(p, s) - \pi(p, s) = \frac{\delta}{1 - \delta} (v - \underline{\pi})$$

⁴⁶If s is continuous and $D(p, s) = \max\{s - p, 0\}$, then $\pi_i(\mathbf{p}, s)$ is linear up to a choke price. As in footnote 12, our analysis applies for this specification, so long as $\underline{s} \geq \bar{s}/2$, so that demand $D(p^m(s), s')$ is non-negative for any monopoly price $p^m(s)$ and demand state s' .

if such a solution exists, and let $p^{\max}(s, v) = p^m(s)$ otherwise. Next, let

$$\pi^{\max}(s, v) = \pi(p^{\max}(s, v), s).$$

Note that $p^{\max}(s, v)$ is non-decreasing in v , and hence—by Assumption 4 and the fact that $p^{\max}(s, v) \leq p^m(s)$ by construction—so is $\pi^{\max}(s, v)$. Finally, let v^* be the greatest fixed point of the equation

$$v = \max_{G \in MPC(F)} \mathbb{E}^G [\min \{\pi^m(x), \pi^{\max}(x, v)\}], \quad (12)$$

which is well-defined by Tarski’s theorem, as $\pi^{\max}(x, v)$ is non-decreasing in v and bounded.

Lemma 1 extends as follows.

Lemma 5 *Optimal per-firm collusive profit equals v^* and is attained by a symmetric, stationary, grim trigger equilibrium. Moreover, a disclosure policy G is optimal if and only if it solves the maximization problem in (12) with $v = v^*$.*

Proof. The construction of a symmetric, stationary, grim trigger equilibrium that attains per-firm collusive profit v^* is the same as in the proof of Lemma 1.

We now show that no equilibrium can attain higher profit. Fix any equilibrium, and let \bar{v} be the supremum over periods t and histories of play up to and including period t of expected per-firm collusive profits from period $t + 1$ onward. Now fix an arbitrary period t and a history of play up to period t , and suppose that when the realized mean demand in period t at this history is s , the equilibrium price vector is $\mathbf{p}(s)$ and firm i ’s equilibrium continuation payoff is v_i . The resulting incentive constraint for firm i is

$$(1 - \delta) \pi_i^d(\mathbf{p}(s), s) + \delta \underline{\pi} \leq (1 - \delta) \pi_i(\mathbf{p}(s), s) + \delta v_i.$$

Averaging this inequality over the n firms, we have

$$\begin{aligned} (1 - \delta) \frac{1}{n} \sum_i \pi_i^d(\mathbf{p}(s), s) &\leq (1 - \delta) \frac{1}{n} \sum_i \pi_i(\mathbf{p}(s), s) + \delta \frac{1}{n} \sum_i (v_i - \underline{\pi}) \\ &\leq (1 - \delta) \frac{1}{n} \sum_i \pi(\mathbf{p}(s), s) + \delta (\bar{v} - \underline{\pi}), \end{aligned}$$

where the second inequality is by definition of \bar{v} , and therefore

$$\frac{1}{n} \sum_i (\pi_i^d(\mathbf{p}(s), s) - \pi_i(\mathbf{p}(s), s)) \leq \frac{\delta}{1-\delta} (\bar{v} - \underline{\pi}).$$

Now, by Assumption 3, there exists $p(s) \geq 0$ such that

$$\begin{aligned} \pi(p(s), s) &\geq \frac{1}{n} \sum_i \pi_i(\mathbf{p}(s), s) \quad \text{and} \\ \pi^d(p(s), s) - \pi(p(s), s) &\leq \frac{\delta}{1-\delta} (\bar{v} - \underline{\pi}). \end{aligned}$$

Moreover, by Assumption 4, $p^{\max}(s, \bar{v})$ maximizes $\pi(p, s)$ subject to $\pi^d(p, s) - \pi(p, s) \leq (\delta/(1-\delta))(\bar{v} - \underline{\pi})$. Therefore, expected collusive profit in period t is at most $\max_{G \in MPC(F)} \mathbb{E}^G [\min \{\pi^m(s), \pi^{\max}(s, \bar{v})\}]$. Since this holds for every period t , we have $\bar{v} \leq \max_{G \in MPC(F)} \mathbb{E}^G [\min \{\pi^m(s), \pi^{\max}(s, \bar{v})\}]$. Hence, $\bar{v} \leq v^*$, by definition of v^* . ■

It remains to solve the information design problem in (12). In general, this problem can be solved using techniques from the static information design literature (e.g., Dworczak and Martini, 2019), and the solution depends on the shape of $\pi^{\max}(s, v)$ as a function of s . However, an explicit solution is available under the following condition, which holds in our leading examples.⁴⁷

Condition 1 $\pi^{\max}(s, v)$ is convex in s for all v .

Under Condition 1, we can define (\hat{s}_L, \hat{s}_H) exactly as in Section 6.2. (For example, with undifferentiated goods, we have $\hat{s}_L = \hat{s}$ as defined in Section 3 and $\hat{s}_H = \bar{s}$.) The next theorem is the general version of our main result (for the stochastic demand case), which generalizes both Theorem 1 and Theorem 3.

Theorem 4 *Assume that Assumptions 1–4 and Condition 1 hold. Then the unique optimal disclosure policy discloses demand states below \hat{s}_L and above \hat{s}_H and conceals demand states*

⁴⁷It suffices that $\pi^{\max}(s, v^*)$ is convex in s , but this weaker condition depends on the endogenous object v^* .

in the interval $[\hat{s}_L, \hat{s}_H]$. The unique optimal collusive price $p(s)$ in state s is given by

$$p(s) = \begin{cases} p^m(s) & \text{if } s < \hat{s}_L, \\ p^m(s^*) & \text{if } s \in [\hat{s}_L, \hat{s}_H], \\ p^{\max}(s, v^*) & \text{if } s > \hat{s}_H. \end{cases}$$

Proof. Analogous to the proof of Theorem 1. Define an auxiliary objective function

$$\tilde{\pi}(s) = \begin{cases} \pi^m(s) & \text{if } s < \hat{s}_L, \\ \frac{\hat{s}_H - s}{\hat{s}_H - \hat{s}_L} \pi^m(\hat{s}_L) + \frac{s - \hat{s}_L}{\hat{s}_H - \hat{s}_L} \pi^{\max}(\hat{s}_H, v^*) & \text{if } s \in [\hat{s}_L, \hat{s}_H], \\ \pi^{\max}(s, v^*) & \text{if } s > \hat{s}_H. \end{cases}$$

Note that $\tilde{\pi}(s)$ is convex and $\tilde{\pi}(s) \geq \min\{\pi^m(s), \pi^{\max}(s, v^*)\}$ for all s . (See Figure 2. In particular, any kink in the objective function $\min\{\pi^m(s), \pi^{\max}(s, v^*)\}$ must be a concave kink, as the functions $\pi^m(s)$ and $\pi^{\max}(s, v^*)$ differ only in that the latter involves an additional constraint for $s \geq s^*$.) Consider the auxiliary problem, $\max_{G \in MPC(F)} \mathbb{E}^G[\tilde{\pi}(s)]$. Since $\tilde{\pi}(s)$ is convex, the solution is full disclosure ($G = F$), and the resulting value is

$$\begin{aligned} & \mathbb{E}^F[\tilde{\pi}(s)] \\ = & F(\hat{s}_L) \mathbb{E}[\pi^m(s) | s < \hat{s}_L] \\ & + (F(\hat{s}_H) - F(\hat{s}_L)) \mathbb{E}\left[\frac{\hat{s}_H - s}{\hat{s}_H - \hat{s}_L} \pi^m(\hat{s}_L) + \frac{s - \hat{s}_L}{\hat{s}_H - \hat{s}_L} \pi^{\max}(\hat{s}_H, v^*) | s \in [\hat{s}_L, \hat{s}_H]\right] \\ & + (1 - F(\hat{s}_H)) \mathbb{E}[\pi^{\max}(s, v^*) | s > \hat{s}_H] \\ = & F(\hat{s}) \mathbb{E}[\pi^m(s) | s \leq \hat{s}_L] + (F(\hat{s}_H) - F(\hat{s}_L)) \pi^m(s^*) + (1 - F(\hat{s}_H)) \mathbb{E}[\pi^{\max}(s, v^*) | s > \hat{s}_H]. \end{aligned}$$

Since $\tilde{\pi}(s) \geq \min\{\pi^m(s), \pi^{\max}(s, v^*)\}$ for all s , this is an upper bound for $\max_{G \in MPC(F)} \mathbb{E}^G[\min\{\pi^m(s), \pi^{\max}(s, v^*)\}]$. But it is attained by disclosing demand states below \hat{s}_L and above \hat{s}_H and concealing demand states in the interval $[\hat{s}_L, \hat{s}_H]$, so this policy is optimal. Moreover, this policy is the unique one that induces only posteriors s where $\tilde{\pi}(s) = \min\{\pi^m(s), \pi^{\max}(s, v^*)\}$, so it is the unique optimal policy. Finally, this disclosure policy is optimal only in conjunction with the prescribed prices. ■

We illustrate Theorem 4 by applying it to the setting of undifferentiated goods with

capacity constraints, in the case where $D(p, s) = s - p$. Assume that the capacity constraint is slack on path but binds off path: C is greater than $\bar{s}/(2n)$, the per-firm monopoly quantity at the highest demand state, but smaller than $\underline{s}/2$, the industry quantity at the lowest demand state with monopoly pricing. Then $p^{\max}(s, v)$ is the larger solution to

$$pC - \frac{p(s-p)}{n} = \frac{\delta}{1-\delta}v,$$

which, after solving the quadratic, gives

$$\pi^{\max}(s, v) = \frac{1}{2} \left(C \sqrt{(nC - s)^2 + 4n \frac{\delta}{1-\delta} v} - C(nC - s) \right) - \frac{\delta}{1-\delta}v.$$

This expression is convex in s , so Condition 1 holds. Hence, in general, a disclose-conceal-disclose disclosure policy is optimal, similarly to the differentiated goods setting in Section 6.2. The logic is again that the objective function $\min\{\pi^m(s), \pi^{\max}(s, v^*)\}$ is piecewise-convex with a concave kink, so it is optimal to disclose the lowest and highest demand states while concealing an intermediate interval of states. We also note that $p^{\max}(s, v)$ and π^{\max} are decreasing in C when $p^{\max}(s, v) < p^m(s)$. Intuitively, reducing capacity decreases firms' deviation payoffs, which facilitates collusion.

B Consumer or Total Surplus Objective

In this appendix, we find the optimal disclosure policy for maximizing a weighted average of producer and consumer surplus. We assume as in Section 4.2 that consumer surplus is quasi-linear in s , so that $CS(p, s) = f(s) + g(p)s + h(p)$ for functions f, g, h . Since $\mathbb{E}[f(s)]$ is independent of the disclosure policy, we consider the problem of maximizing a weighted average of $\mathbb{E}[g(p)s + h(p)]$ and $\mathbb{E}[\Pi(p, s)]$, with weights $1 - \alpha$ and α , for $\alpha \in [0, 1)$ (the $\alpha = 1$ case already having been considered in the text). We assume that, for any disclosure policy G , the firm-optimal subgame perfect equilibrium is played, resulting in an expected profit V given by the greatest solution to

$$V = \mathbb{E}^G [\min \{\Pi^m(x), \Pi^{\max}(n, \delta, V)\}],$$

and a price $p(x, V)$ at disclosed mean demand state x given by $p(x, V) = p^m(x)$ if $p^m(x) \leq \Pi^{\max}(n, \delta, V)$, and $p(x, V) = \min\{p : \Pi(x, s) = \Pi^{\max}(n, \delta, V)\}$ otherwise. Given this function $p(x, V)$, the designer's problem is to first solve the "inner problem,"

$$\begin{aligned} & \max_{G \in MPC(F)} \mathbb{E}^G [(1 - \alpha)(g(p(x, V))x + h(p(x, V))) + \alpha \min\{\Pi^m(x), \Pi^{\max}(n, \delta, V)\}] \\ \text{s.t.} \quad & V \text{ is the greatest solution to } V = \mathbb{E}^G [\min\{\Pi^m(x), \Pi^{\max}(n, \delta, V)\}], \end{aligned}$$

for each V that is the greatest solution to $V = \mathbb{E}^G [\min\{\Pi^m(x), \Pi^{\max}(n, \delta, V)\}]$ for some $G \in MPC(F)$, and then to maximize over V .

The inner problem can be rewritten as an unconstrained information design problem by letting β be the multiplier on the constraint (which exists by Slater's condition for any $V < V^*$, the solution when $\alpha = 1$), normalizing the objective by $1/(1 - \alpha)$, and letting $\lambda = (\alpha + \beta)/(1 - \alpha)$, to obtain

$$\max_{G \in MPC(F)} \mathbb{E}^G [(g(p(x, V))x + h(p(x, V))) + \lambda \min\{\Pi^m(x), \Pi^{\max}(n, \delta, V)\}]. \quad (13)$$

In general, (13) can be solved by standard techniques (e.g., Dworczak and Martini, 2019), given the functions p , g , h , and Π^m and the multiplier λ . Here, we consider the leading case of linear demand with an unknown intercept or an unknown constant marginal cost.

Unknown demand. As shown in Section 4.2, $\Pi(p, x) = p(x - p)$, $g(p) = -p$, and $h(p) = p^2/2$. We thus have

$$\Pi^m(x) = \frac{x^2}{4} \quad \text{and} \quad g(p^m(x))x + h(p^m(x)) = -\frac{3x^2}{8}.$$

Suppressing the argument of $\Pi^{\max}(n, \delta, V)$ and letting $p^{\max}(x, V)$ solve $\Pi(p^{\max}(x, V), x) = \Pi^{\max}$ (for $x \geq s^* \equiv 2\sqrt{\Pi^{\max}}$), we calculate

$$\begin{aligned} p^{\max}(x, V) &= \frac{x - \sqrt{x^2 - 4\Pi^{\max}}}{2} \quad \text{and} \\ g(p^{\max}(x, V))x + h(p^{\max}(x, V)) &= \frac{-x^2 - 2\Pi^{\max} + x\sqrt{x^2 - 4\Pi^{\max}}}{4}. \end{aligned}$$

The designer's (inner) problem is thus

$$\max_{G \in MPC(F)} \mathbb{E}^G \left[\begin{array}{c} \mathbf{1}\{x < s^*\} (-3 + 2\lambda) \frac{x^2}{8} \\ + \mathbf{1}\{x \geq s^*\} \left(\frac{-x^2 + x\sqrt{x^2 - 4\Pi^{\max}}}{4} + \left(-\frac{1}{2} + \lambda \right) \Pi^{\max} \right) \end{array} \right].$$

Note that if $\lambda < 3/2$ then the objective is decreasing and concave in x for $x < s^*$ and is increasing and concave in x for $x \geq s^*$ (as the function $-x^2 + x\sqrt{x^2 - 4\Pi^{\max}}$ is increasing and concave for $x \geq 2\sqrt{\Pi^{\max}}$), with a convex kink at s^* . Thus, if the optimal multiplier λ (i.e., the multiplier in the inner problem with the optimal value for V) is less than $3/2$, the optimal policy is a binary signal that discloses only whether demand is above or below a cutoff s^{**} , by an argument similar to the proof of Proposition 4 of Dworczak and Martini (2019). Intuitively, when $\lambda < 3/2$ the objective has the same shape as consumer surplus $g(p(x))x + h(p(x))$; and when $x < s^*$, (monopoly) price is linearly increasing in x , so consumer surplus is decreasing and concave in x ; whereas when $x \geq s^*$, price is decreasing and convex in x (by Rotemberg and Saloner's logic), so consumer surplus is increasing and concave in x . The optimal policy is thus a binary signal that discloses only whether demand is "low" (in which case firms set the corresponding monopoly price) or "high" (causing a "price war").

If instead the optimal multiplier λ is greater than $3/2$, the objective is increasing and convex for $x < s^*$, so the objective is S-shaped overall. In this case, the solution is upper censorship with a cutoff $\hat{s} < s^*$, as in the problem considered in the text.

The next result summarizes this discussion and also shows that a binary signal maximizes consumer surplus. (Equivalently, the optimal multiplier is less than $3/2$ when $\alpha = 0$.) This follows because the optimal signal for any weight α is either a binary signal or is upper censorship with a cutoff $\hat{s} \in (s, s^*)$, but given upper censorship with a cutoff \hat{s} , pooling all states below \hat{s} in a single signal realization increases expected consumer surplus, so the optimal signal when $\alpha = 0$ must be binary.

Proposition 4 *Under linear demand with an unknown intercept, the optimal disclosure policy is either a binary signal that reveals only whether demand is below or above a cutoff; or it is upper censorship. The former policy is optimal if the weight on consumer surplus is*

sufficiently high; the latter policy is optimal if the weight on producer surplus is sufficiently high.

Proof. Given the above discussion, it suffices to show that the optimal multiplier λ^* is at most $3/2$ when $\alpha = 0$. To see this, suppose toward a contradiction that $\lambda^* > 3/2$. Then, as shown above, the optimal disclosure policy is an upper censorship policy G with a cutoff $\hat{s} < s^*$. We can also assume that $\hat{s} > \underline{s}$, as otherwise this policy is no disclosure, which is also a binary signal.

To obtain a contradiction, we show that the binary signal \hat{G} that discloses only whether demand is above or below \hat{s} yields strictly higher expected consumer surplus than G . To see this, let V and \hat{V} denote expected profit under G and \hat{G} , respectively. Recall that $p(x, V) = p^m(x)$ for all $x < \hat{s}$ (as $\hat{s} < s^*$) and that $g(p^m(x))x + h(p^m(x))$ is strictly concave in x , so \hat{G} yields strictly higher expected consumer surplus than G if $p(\mathbb{E}^F[s|s < \hat{s}], \hat{V}) \leq p(\mathbb{E}^F[s|s < \hat{s}], V)$ and $p(\mathbb{E}^F[s|s \geq \hat{s}], \hat{V}) \leq p(\mathbb{E}^F[s|s \geq \hat{s}], V)$. In turn, since $p(x, \hat{V})$ is increasing in \hat{V} , it suffices to show that $V \geq \hat{V}$. But this holds because V and \hat{V} are, respectively, the greatest fixed points of the equations

$$\tilde{V} = \mathbb{E}^G \left[\min \left\{ \Pi^m(x), \Pi^{\max}(\delta, n, \tilde{V}) \right\} \right] \quad \text{and} \quad \tilde{V} = \mathbb{E}^{\hat{G}} \left[\min \left\{ \Pi^m(x), \Pi^{\max}(\delta, n, \tilde{V}) \right\} \right];$$

but, for any \tilde{V} , we have

$$\begin{aligned} \mathbb{E}^G \left[\min \left\{ \Pi^m(x), \Pi^{\max}(\delta, n, \tilde{V}) \right\} \right] &= \int_{\underline{s}}^{\hat{s}} \Pi^m(s) dF(s) + (1 - F(\hat{s})) \Pi^{\max}(\delta, n, \tilde{V}) \\ &\geq F(\hat{s}) \Pi^m(\mathbb{E}^F[s|s < \hat{s}]) + (1 - F(\hat{s})) \Pi^{\max}(\delta, n, \tilde{V}) \\ &\geq F(\hat{s}) \min \left\{ \Pi^m(\mathbb{E}^F[s|s < \hat{s}]), \Pi^{\max}(\delta, n, \tilde{V}) \right\} \\ &\quad + (1 - F(\hat{s})) \min \left\{ \Pi^m(\mathbb{E}^F[s|s \geq \hat{s}]), \Pi^{\max}(\delta, n, \tilde{V}) \right\} \\ &= \mathbb{E}^{\hat{G}} \left[\min \left\{ \Pi^m(x), \Pi^{\max}(\delta, n, \tilde{V}) \right\} \right], \end{aligned}$$

where the first inequality is by convexity of $\Pi^m(s)$, so the greatest fixed point of the first equation is not lower than that of the second. ■

A natural conjecture is that there exists a weight on producer surplus $\alpha^* \in (0, 1)$ such

that if $\alpha < \alpha^*$ then the optimal multiplier λ is less than $3/2$, so a binary signal is optimal; and if $\alpha \geq \alpha^*$ then $\lambda \geq 3/2$, so upper censorship is optimal. This holds whenever the optimal multiplier β on the constraint $\mathbb{E}^G [\Pi^m(x), \Pi^{\max}(n, \delta, V)] = V$ is monotone in α , as then λ is monotone in α , so that $\lambda \geq 3/2$ iff α exceeds a cutoff α^* . However, we do not have a proof that β is monotone in α .

Unknown cost. Since $\Pi(p, x) = (p - x)(1 - p)$ and $CS(p, x) = h(p)$ for $h(p) = (1/2)(1 - p)^2$, we have

$$p^m(x) = \frac{1+x}{2}, \quad \Pi^m(x) = \frac{(1-x)^2}{4}, \quad \text{and} \quad h(p^m(x)) = \frac{(1-x)^2}{8}.$$

Suppressing the argument of $\Pi^{\max}(n, \delta, V)$ and letting $p^{\max}(x, V)$ solve $\Pi(p^{\max}(x, V), x) = \Pi^{\max}$ (for $x \leq s^* \equiv 1 - 2\sqrt{\Pi^{\max}}$), we calculate

$$\begin{aligned} p^{\max}(x, V) &= \frac{1+x - \sqrt{(1-x)^2 - 4\Pi^{\max}}}{2} \quad \text{and} \\ h(p^{\max}(x, V)) &= \frac{(1-x)^2 - 2\Pi^{\max} + (1-x)\sqrt{(1-x)^2 - 4\Pi^{\max}}}{4}. \end{aligned}$$

The designer's (inner) problem is thus

$$\max_{G \in MPC(F)} \mathbb{E}^G \left[\mathbf{1}\{x \leq s^*\} \left(\frac{(1-x)^2 + (1-x)\sqrt{(1-x)^2 - 4\Pi^{\max}}}{4} + \left(-\frac{1}{2} + \lambda\right) \Pi^{\max} \right) \right. \\ \left. + \mathbf{1}\{x > s^*\} (1+2\lambda) \frac{(1-x)^2}{8} \right].$$

Note that the objective is decreasing and convex for $x \leq s^*$ (as the function $(1-x)^2 + (1-x)\sqrt{(1-x)^2 - 4\Pi^{\max}}$ is decreasing and convex for $x \leq s^*$); whereas for $x > s^*$, the objective is increasing and concave if $\lambda \leq -1/2$ and is decreasing and convex if $\lambda > -1/2$. In the former case, the objective consists of a decreasing and convex piece followed by an increasing and concave piece, and it is straightforward to show that the optimal disclosure policy is upper censorship. In the latter case, the objective consists of two decreasing and convex pieces that meet at a kink, and in general the optimal disclosure policy censors an intermediate range of states around the kink and discloses the lowest and highest states.

However, unlike in the unknown demand case, we have not been able to determine which of these cases applies for the problem of maximizing consumer surplus ($\alpha = 0$).⁴⁸

C Proof of Proposition 2.3

Fix any convex function $\Pi^m(s)$ and any constant Π^{\max} . Let (F_1, F_2, G_1, G_2) be such that

$$\begin{aligned} F_1 &\in MPC(F_2), \\ G_1 &\in \underset{G \in MPC(F_1)}{\operatorname{argmax}} \mathbb{E}^G [\min \{\Pi^m(x), \Pi^{\max}\}], \quad \text{and} \\ G_2 &\in \underset{G \in MPC(F_2)}{\operatorname{argmax}} \mathbb{E}^G [\min \{\Pi^m(x), \Pi^{\max}\}]. \end{aligned}$$

We show that $G_1 \in MPC(G_2)$, or equivalently $\int_{\underline{s}}^s G_1(s) ds \leq \int_{\underline{s}}^s G_2(s) ds$ for all x (since G_1 and G_2 have the same mean). As shown in the text, this completes the proof of Proposition 2.3.

By Theorem 1, we have

$$G_1(s) = \begin{cases} F_1(s) & \text{if } s \leq \hat{s}_1, \\ F_1(\hat{s}_1) & \text{if } \hat{s}_1 < s < s^*, \\ 1 & \text{if } s \geq s^*, \end{cases}$$

where $s^* = E^{F_1}[s | s > \hat{s}_1]$ satisfies $\Pi^m(s^*) = \Pi^{\max}$, and similarly for G_2 . Note that, for any $s \leq \hat{s}_2$, we have

$$\int_{\underline{s}}^s (G_1(s) - G_2(s)) ds = \int_{\underline{s}}^s (G_1(s) - F_2(s)) ds \leq \int_{\underline{s}}^s (F_1(s) - F_2(s)) ds \leq 0,$$

where the equality is by $G_2(s) = F_2(s)$ for all $s \leq \hat{s}_2$, the first inequality is by $G_1(s) \leq F_1(s)$ for all $s < s^*$, and the second inequality is because $F_1 \in MPC(F_2)$. Next, $G_1(s) - G_2(s)$ is non-decreasing on the interval $[\hat{s}_2, s^*)$, as on this interval $G_1(s)$ is non-decreasing and $G_2(s) = F_2(\hat{s}_2)$ is constant. In addition, $G_1(s) - G_2(s) = 0$ for $s \geq s^*$. Thus,

⁴⁸Conversely, as $\lambda \rightarrow \infty$ (approaching the $\alpha = 1$ case considered in the text), the objective flattens out for $s \leq s^*$, implying that lower censorship is optimal for sufficiently high λ .

$\int_{\underline{s}}^s (G_1(s) - G_2(s)) ds$ is convex on $[\hat{s}_2, s^*]$ and constant on $(s^*, \bar{s}]$. Therefore, since $\int_{\underline{s}}^s (G_1(s) - G_2(s)) ds \leq 0$ for all $s \leq \hat{s}_2$, if $\int_{\underline{s}}^s (G_1(s) - G_2(s)) ds$ is ever strictly positive then for some $s \in (s^*, \bar{s}]$, then it must be strictly positive at $s = s^*$ (since, as a convex function on $[\hat{s}_2, s^*]$, it is bounded above by its linear interpolation over this interval). But, by integration by parts,

$$\begin{aligned}\mathbb{E}^{G_1}[s] &= \bar{s} - \int_{\underline{s}}^{\bar{s}} G_1(s) ds = s^* - \int_{\underline{s}}^{s^*} G_1(s) ds, \quad \text{and} \\ \mathbb{E}^{G_2}[s] &= \bar{s} - \int_{\underline{s}}^{\bar{s}} G_2(s) ds = s^* - \int_{\underline{s}}^{s^*} G_2(s) ds,\end{aligned}$$

so since $\mathbb{E}^{G_1}[s] = \mathbb{E}^{G_2}[s]$ we have $\int_{\underline{s}}^{s^*} (G_1(s) - G_2(s)) ds = 0$. Thus, $\int_{\underline{s}}^s (G_1(s) - G_2(s)) ds \leq 0$ for all s , completing the proof.

D Example with Persistent Demand

Assume a binary demand state ($s \in \{\underline{s}, \bar{s}\}$) and linear demand ($\Pi(p, s) = p(s - p)$), and consider the parameters $\underline{s} = 1$, $\bar{s} = 2$, $\delta = .55$, $n = 2$, and $\Pr(s_{t+1} = \bar{s} | s_t = \bar{s}) = \Pr(s_{t+1} = \underline{s} | s_t = \underline{s}) = \rho \in (1/2, 1)$. Binary demand violates our assumption that the distribution of states is atomless; however, that assumption is easily relaxed. With a binary state, an upper censorship policy now corresponds to disclosing state \underline{s} with some probability q (conditional on $s = \underline{s}$) and pooling state \underline{s} together with state \bar{s} otherwise. Upper censorship is optimal by essentially the same proof as in the atomless case, so we can parameterize an optimal disclosure policy by $q \in [0, 1]$, with $q = 0$ being no disclosure, $q \in (0, 1)$ being non-trivial upper censorship, and $q = 1$ being full disclosure.

Figures 3–5 display the optimal disclosure policy, firm profit, and consumer surplus as ρ ranges from $1/2$ to 1 . In Figure 3, the blue curve plots the optimal disclosure policy q at last-period demand state \underline{s} , while the orange line plots q at last-period demand state \bar{s} . Note that the blue curve is always above the orange curve, as under positive persistence ($\rho > 1/2$) more information is disclosed when the last-period state is lower, as shown in Theorem 2. In addition, Figure 3 displays three distinct equilibrium regimes. In Regime 1, non-trivial upper censorship is optimal for both last-period demand states. For $\rho \in (1/2, 0.727)$, Regime

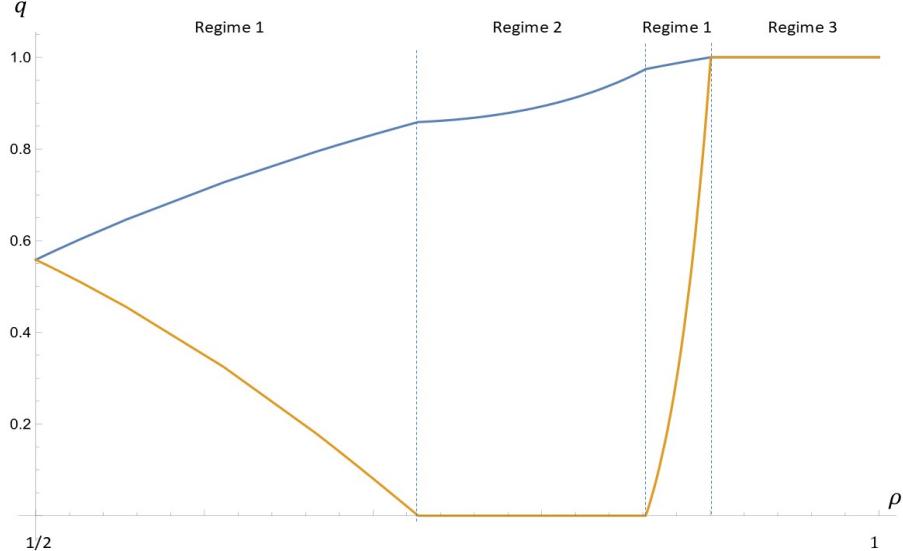


Figure 3: Optimal disclosure policy. The blue curve is the probability of disclosing \underline{s} when the current state is \underline{s} and the last-period state was \underline{s} ; the orange curve is the probability of disclosing \underline{s} when the current state is \underline{s} and the last-period state was \bar{s} .

1 prevails, and increasing ρ leads to more disclosure at last-period state \underline{s} and less disclosure at last-period state \bar{s} . Intuitively, increasing ρ makes firms more pessimistic at last-period state \underline{s} and more optimistic at last-period state \bar{s} , which increases disclosure at last-period state \underline{s} and decreases disclosure at last-period state \bar{s} . Once ρ reaches 0.727, firms are so optimistic at last-period state \bar{s} that no disclosure becomes optimal, while further increases in ρ continue to increase disclosure at last-period state \underline{s} . This second regime persists until ρ reaches 0.863. At this point, demand is so persistent that future profits are much higher at last-period state \bar{s} than at last-period state \underline{s} , which makes partial disclosure optimal again at last-period state \bar{s} , so the equilibrium is again in Regime 1. Further increases in ρ then rapidly increase disclosure for both last-period states, until ρ reaches 0.902, at which point full disclosure becomes optimal for both last-period states.

Figures 4 and 5 trace the implications of these effects for firm profit and consumer surplus. In Figure 4, the blue curve plots a firm's continuation value (discounted sum of profits) at last-period demand state \underline{s} ; the orange curve plots this value at last-period demand state \bar{s} ; and the green curve is the average of the two, which equals a firm's ex ante expected profit. In Regime 1, increasing ρ decreases the continuation value at last-period state \underline{s} and increases it at last-period state \bar{s} . The net effect is to (slightly) increase expected profits, as increasing

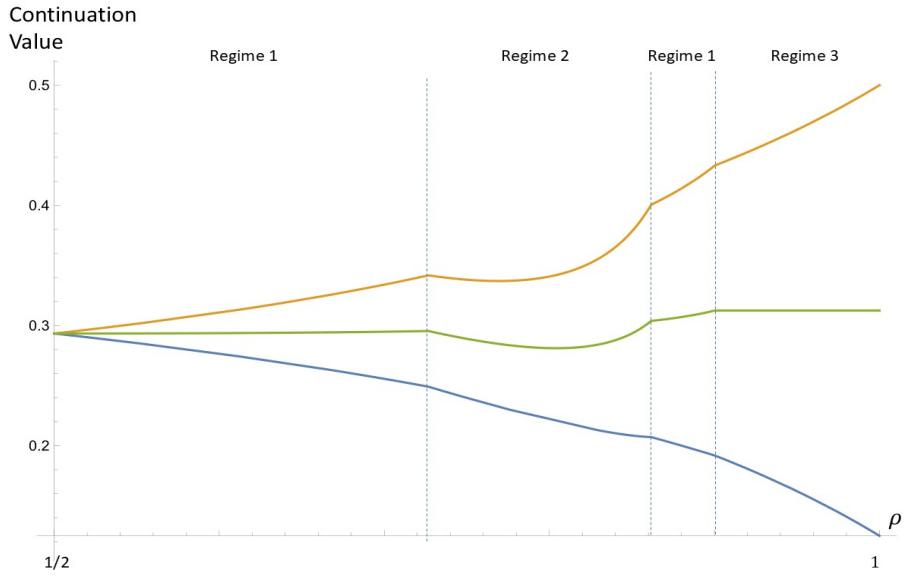


Figure 4: Optimal continuation values. The blue curve is a firm's continuation value at last-period state \underline{s} . The orange curve is the corresponding value at last-period state \bar{s} . The green curve is the ex ante expected profit.

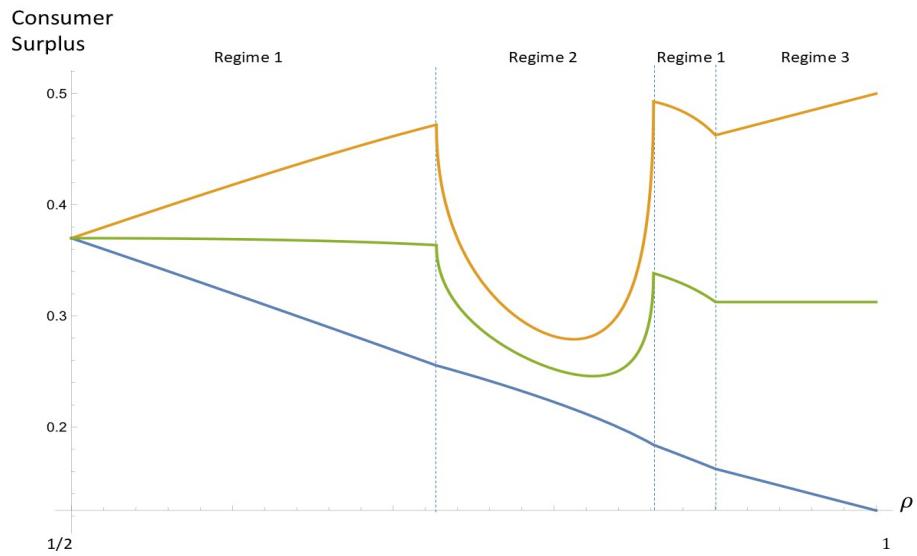


Figure 5: Consumer surplus. The blue curve is the current-period consumer surplus at last-period state \underline{s} . The orange line is the corresponding value at last-period state \bar{s} . The green curve is ex ante expected consumer surplus.

the continuation value at last-period state \bar{s} relaxes the binding incentive constraint. In contrast, the effect of increasing ρ on profits is non-monotone in Regime 2 and is zero in Regime 3 (where optimal profits are first-best). In Figure 5, the blue curve plots the current-period consumer surplus at last-period demand state \underline{s} ; the orange curve plots it at last-period demand state \bar{s} ; and the green curve is the average of the two, which equals ex ante expected consumer surplus. Expected consumer surplus is decreasing in ρ in Regime 1 (albeit only slightly when $\rho \in (1/2, 0.727)$), non-monotone in ρ in Regime 2, and constant in ρ in Regime 3.

References

- [1] Abreu, Dilip. “Extremal Equilibria of Oligopolistic Supergames.” *Journal of Economic Theory* 39.1 (1986): 191-225.
- [2] Abreu, Dilip. “On the Theory of Infinitely Repeated Games with Discounting.” *Econometrica* 56.2 (1988): 383-396.
- [3] Abreu, Dilip, Paul Milgrom, and David Pearce. “Information and Timing in Repeated Partnerships.” *Econometrica* 59.6 (1991): 1713-1733.
- [4] Amir, Rabah, Philip Erickson, and Jim Jin. “On the Microeconomic Foundations of Linear Demand for Differentiated Products.” *Journal of Economic Theory* 168 (2017): 641-665.
- [5] Asker, John, Chaim Fershtman, and Ariel Pakes. “The Impact of Artificial Intelligence Design on Pricing.” *Journal of Economics & Management Strategy* 33.2 (2024): 276-304.
- [6] Assad, Stephanie, et al. “Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market.” *Journal of Political Economy* 132.3 (2024): 723-771.
- [7] Athey, Susan, and Kyle Bagwell. “Optimal Collusion with Private Information.” *RAND Journal of Economics* 32.3 (2001): 428-465.
- [8] Athey, Susan, Kyle Bagwell, and Chris Sanchirico. “Collusion and Price Rigidity.” *Review of Economic Studies* 71.2 (2004): 317-349.
- [9] Awaya, Yu, and Vijay Krishna. “On Communication and Collusion.” *American Economic Review* 106.2 (2016): 285-315.
- [10] Bagwell, Kyle, and Robert W. Staiger. “Collusion over the Business Cycle.” *Rand Journal of Economics* 28.1 (1997): 82-106.

- [11] Banchio, Martino, and Giacomo Mantegazza. “Artificial Intelligence and Spontaneous Collusion.” *arXiv preprint arXiv:2202.05946* (2024).
- [12] Bernheim, B. Douglas, and Erik Madsen. “Price Cutting and Business Stealing in Imperfect Cartels.” *American Economic Review* 107.2 (2017): 387-424.
- [13] Blackwell, David. “Equivalent Comparisons of Experiments.” *Annals of Mathematical Statistics* (1953): 265-272.
- [14] Bonatti, Alessandro, Raffaele Ficco, and Salvatore Piccolo. “Data Analytics for Algorithm Pricing.” *Working Paper* (2024).
- [15] Bordoli, Davide, “Selling Information with Strategic Externalities.” *Working Paper* (2025).
- [16] Brown, Zach Y., and Alexander MacKay. “Competition in Pricing Algorithms.” *American Economic Journal: Microeconomics* 15.2 (2023): 109-156.
- [17] Buehler, Stefan, and Dennis L. Gärtner. “Making Sense of Nonbinding Retail-Price Recommendations.” *American Economic Review* 103.1 (2013): 335-359.
- [18] Calder-Wang, Sophie, and Gi Heung Kim. “Coordinated vs Efficient Prices: The Impact of Algorithmic Pricing on Multifamily Rental Markets.” *Working Paper*, University of Pennsylvania, SSRN 4403058 (2023).
- [19] Calvano, Emilio, et al. “Protecting Consumers from Collusive Prices due to AI.” *Science* 370.6520 (2020a): 1040-1042.
- [20] Calvano, Emilio, et al. “Artificial Intelligence, Algorithmic Pricing, and Collusion.” *American Economic Review* 110.10 (2020b): 3267-3297.
- [21] Carlton, Dennis. “The Rigidity of Prices.” *American Economic Review* 76.4 (1986): 637-58.
- [22] Cooper, William L., Tito Homem-de-Mello, and Anton J. Kleywegt. “Learning and Pricing with Models that do not Explicitly Incorporate Competition.” *Operations Research* 63.1 (2015): 86-103.
- [23] Dworczak, Piotr, and Giorgio Martini. “The Simple Economics of Optimal Persuasion.” *Journal of Political Economy* 127.5 (2019): 1993-2048.
- [24] Ellison, Glenn. “Theories of Cartel Stability and the Joint Executive Committee.” *Rand Journal of Economics* 25.1 (1994): 37-57.
- [25] Ely, Jeffrey C. “Beeps.” *American Economic Review* 107.1 (2017): 31-53.
- [26] Ezrachi, Ariel, and Maurice E. Stucke. “Artificial Intelligence & Collusion: When Computers Inhibit Competition.” *University of Illinois Law Review* (2017): 1775.

- [27] Farboodi, Maryam, Nima Haghpanah, and Ali Shourideh. “Good Data and Bad Data: The Welfare Effects of Price Discrimination.” *arXiv preprint arXiv:2502.03641* (2025).
- [28] Gentzkow, Matthew, and Emir Kamenica. “A Rothschild-Stiglitz Approach to Bayesian Persuasion.” *American Economic Review* 106.5 (2016): 597-601.
- [29] Gerlach, Heiko. “Stochastic Market Sharing, Partial Communication and Collusion.” *International Journal of Industrial Organization* 27.6 (2009): 655-666.
- [30] Goltsman, Maria, and Gregory Pavlov. “Communication in Cournot oligopoly.” *Journal of Economic Theory* 153 (2014): 152-176.
- [31] Green, Edward J., and Robert H. Porter. “Noncooperative Collusion under Imperfect Price Information.” *Econometrica* 52.1 (1984): 87-100.
- [32] Haltiwanger, John, and Joseph E. Harrington Jr. “The Impact of Cyclical Demand Movements on Collusive Behavior.” *RAND Journal of Economics* 22.1 (1991): 89-106.
- [33] Hanazono, Makoto, and Huanxing Yang. “Collusion, Fluctuating Demand, and Price Rigidity.” *International Economic Review* 48.2 (2007): 483-515.
- [34] Hansen, Karsten T., Kanishka Misra, and Mallesh M. Pai. “Frontiers: Algorithmic Collusion: Supra-Competitive Prices via Independent Algorithms.” *Marketing Science* 40.1 (2021): 1-12.
- [35] Harrington Jr, Joseph E. “How do Cartels Operate?.” *Foundations and Trends in Microeconomics* 2.1 (2006): 1-105.
- [36] Harrington Jr, Joseph E. “Detecting Cartels.” *Handbook of Antitrust Economics*, Buccirossi ed. (2008).
- [37] Harrington Jr, Joseph E. “The Effect of Outsourcing Pricing Algorithms on Market Competition.” *Management Science* 68.9 (2022): 6889-6906.
- [38] Harrington Jr, Joseph E. “Hub-and-Spoke Collusion with a Third-Party Pricing Algorithm.” *Working Paper* (2025).
- [39] Harrington, Joseph E., and Juan Ortner. “The Equilibrium Design of a Third-Party’s Pricing Algorithm.” *Working Paper* (2025).
- [40] Harrington, Joseph E., and Andrzej Skrzypacz. “Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices.” *American Economic Review* 101.6 (2011): 2425-2449.
- [41] Hickok, Nathaniel. “Algorithm Design Meets Information Design: Price Recommendation Algorithms on Online Platforms.” *Working Paper*, MIT.
- [42] Kamenica, Emir, and Matthew Gentzkow. “Bayesian Persuasion.” *American Economic Review* 101.6 (2011): 2590-2615.

- [43] Kandori, Michihiro. "The Use of Information in Repeated Games with Imperfect Monitoring." *Review of Economic Studies* 59.3 (1992): 581-593.
- [44] Kandori, Michihiro. "Correlated Demand Shocks and Price Wars during Booms." *Review of Economic Studies* 58.1 (1991): 171-180.
- [45] Kawai, Kei, Jun Nakabayashi, and Juan M. Ortner. "The Value of Privacy in Cartels: An Analysis of the Inner Workings of a Bidding Ring." *Review of Economic Studies* (2025): Forthcoming.
- [46] Kolotilin, Anton. "Optimal Information Disclosure: A Linear Programming Approach." *Theoretical Economics* 13.2 (2018): 607-635.
- [47] Kolotilin, Anton, et al. "Persuasion of a Privately Informed Receiver." *Econometrica* 85.6 (2017): 1949-1964.
- [48] Kolotilin, Anton, and Hongyi Li. "Relational Communication." *Theoretical Economics* 16.4 (2021): 1391-1430.
- [49] Kolotilin, Anton, Timofiy Mylovanov, and Andriy Zapechelnyuk. "Censorship as Optimal Persuasion." *Theoretical Economics* 17.2 (2022): 561-585.
- [50] Kolotilin, Anton, Roberto Corrao, and Alexander Wolitzky. "Persuasion and Matching: Optimal Productive Transport." *Journal of Political Economy* (2024): Forthcoming.
- [51] Kolotilin, Anton, and Alexander Wolitzky. "The Economics of Partisan Gerrymandering." *Econometrica* (2025), Forthcoming.
- [52] Klein, Timo. "Autonomous Algorithmic Collusion: Q-Learning under Sequential Pricing." *RAND Journal of Economics* 52.3 (2021): 538-558.
- [53] Kuvalekar, Aditya, Elliot Lipnowski, and Joao Ramos. "Goodwill in Communication." *Journal of Economic Theory* 203 (2022): 105467.
- [54] Lamba, Rohit, and Sergey Zhuk. "Pricing with Algorithms." *arXiv preprint arXiv:2205.04661* (2024).
- [55] Marshall, Robert C., and Leslie M. Marx. *The Economics of Collusion: Cartels and Bidding Rings*. Mit Press, 2012.
- [56] Martin, Simon, and Alexander Rasch. "Demand Forecasting, Signal Precision, and Collusion with Hidden Actions." *International Journal of Industrial Organization* 92 (2024): 103036.
- [57] Maskin, Eric, and Jean Tirole. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles." *Econometrica* 56.3 (1988): 571-599.
- [58] McAfee, R. Preston, and John McMillan. "Bidding Rings." *American Economic Review* (1992): 579-599.

[59] Mehra, Salil K. "Antitrust and the robo-seller: Competition in the Time of Algorithms." *Minnesota Law Review* 100 (2015): 1323.

[60] Miklós-Thal, Jeanine, and Catherine Tucker. "Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers?." *Management Science* 65.4 (2019): 1552-1561.

[61] O'Connor, Jason, and Nathan E. Wilson. "Reduced Demand Uncertainty and the Sustainability of Collusion: How AI Could Affect Competition." *Information Economics and Policy* 54 (2021): 100882.

[62] Ortner, Juan, Takuo Sugaya, and Alexander Wolitzky. "Mediated Collusion." *Journal of Political Economy* 132.4 (2024): 1247-1289.

[63] Pigou, A.C. *The Economics of Welfare*. (1920).

[64] Porter, Robert H. "A Study of Cartel Stability: the Joint Executive Committee, 1880-1886." *Bell Journal of Economics* (1983): 301-314.

[65] Plan, Asaf. "Symmetry in n-Player Games." *Journal of Economic Theory* 207 (2023): 105549.

[66] Renault, Jérôme, Eilon Solan, and Nicolas Vieille. "Optimal Dynamic Information Provision." *Games and Economic Behavior* 104 (2017): 329-349.

[67] Rotemberg, Julio J., and Garth Saloner. "A Supergame-Theoretic Model of Price Wars During Booms." *American Economic Review* 76.3 (1986): 390-407.

[68] Sahuguet, Nicolas, and Alexis Walckiers. "A Theory of Hub-and-Spoke Collusion." *International Journal of Industrial Organization* 53 (2017): 353-370.

[69] Salcedo, Bruno. "Pricing Algorithms and Tacit Collusion." *Working Paper*, Pennsylvania State University (2015).

[70] Smolin, Alex and Takuro Yamashita. "Information Design in Smooth Games." *Working Paper*, Toulouse School of Economics (2025).

[71] Skrzypacz, Andrzej, and Hugo Hopenhayn. "Tacit Collusion in Repeated Auctions." *Journal of Economic Theory* 114.1 (2004): 153-169.

[72] Strassen, Volker. "The Existence of Probability Measures with Given Marginals." *Annals of Mathematical Statistics* 36.2 (1965): 423-439.

[73] Sugaya, Takuo, and Alexander Wolitzky. "Maintaining Privacy in Cartels." *Journal of Political Economy* 126.6 (2018): 2569-2607.

[74] US Department of Justice. *United States of America et al. v. RealPage* (2024a).

[75] US Department of Justice. "Justice Department Sues RealPage for Algorithmic Pricing Scheme that Harms Millions of American Renters" (2024b).

- [76] US Department of Justice. *United States of America et al. v. RealPage, Proposed Final Judgement* (2025).
- [77] Vives, Xavier. *Oligopoly Pricing: Old Ideas and New Tools*. MIT Press, Cambridge, Mass., 2001.
- [78] Waltman, Ludo, and Uzay Kaymak. “Q-Learning Agents in a Cournot Oligopoly Model.” *Journal of Economic Dynamics and Control* 32.10 (2008): 3275-3293.
- [79] Wu, Karen. “Welfare Effects of Market Research.” *Working Paper* (2025).