Collusion with Optimal Information Disclosure*

Takuo Sugaya Alexander Wolitzky
Stanford GSB MIT

January 17, 2025

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 optimally discloses demand or cost information to maximize firms' collusive profit. We assume that profit is affine in the unknown state, so expected profit is determined by the expected state. We show that an *upper censorship* disclosure policy is optimal, which leads to *price rigidity* and *supra-monopoly prices* at some states. Under a concavity condition, improving the algorithm's accuracy reduces expected consumer surplus. 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.

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

JEL codes: C73, D43, D82

*For helpful comments, we thank Kyle Bagwell, Alessandro Bonatti, Sophie Calder-Wang, Glenn Ellison, Drew Fudenberg, Joe Harrington, Anton Kolotilin, Stephen Morris, Juan Ortner, 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 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 recently arisen in a number of other industries, including retail gasoline pricing (A2i Systems and Kalibrate; see Assad et al., 2024) and hotel room pricing (Rainmaker; see Harrington, 2024). In addition, related concerns have also been raised regarding some 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 demand or cost information than individual firms can selectively disclose this information to maximize the firms' collusive profit.[1] We work in the canonical setting of repeated Bertrand competition with stochastic demand, introduced by Rotemberg and Saloner (1986).[2] Following Rotemberg and Saloner, our baseline setting

[1]The 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 a superset of the firms'. We discuss these issues later on.

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

assumes undifferentiated 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 the key technical assumption that profit is affine in the unknown state, so that, for 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.

Our main result is that optimal information disclosure takes a simple *upper censorship* form: 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 the demand state 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.

The logic of these results is fairly straightforward. As in Rotemberg and Saloner's model, 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's analysis—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, 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. The technical reason why this is so is that the firms' "capped monopoly profit," $\min\{\Pi^m(s), \Pi^{\max}\}$—where $\Pi^{\max}$ is the maximum industry profit that the firms can attain in equilibrium in a single period without violating incentive constraints, which is independent of s with undifferentiated products and iid demand—is a "convex-then-concave" function of s, and it is well-known that the optimal disclosure policy with a convex-then-concave objective function is upper censorship (Kolotilin, 2018; Dworczak and Martini, 2019; Kolotilin et al., 2022).

The optimal collusive equilibrium displays clean comparative statics. A decrease in the number of firms or an increase in the discount factor increases collusive profit and—more interestingly—makes collusive prices more flexible. The logic is that higher profits and more flexible prices go hand-in-hand, because the purpose of rigid prices is to deter deviations, which is less necessary when equilibrium continuation payoffs are higher. We also show that prices are higher (and hence consumer surplus is lower) at each demand state under the optimal information disclosure policy as compared to the full disclosure case studied by Rotemberg and Saloner. Finally, we provide a simple condition—concavity of consumer surplus in s under monopoly pricing—under which improving the algorithm's accuracy reduces expected consumer surplus. The latter two results speak directly to antitrust concerns surrounding algorithmic information-sharing. Specifically, while prior studies have found an ambiguous affect of improved algorithmic demand prediction on consumer surplus (Sugaya and Wolitzky, 2018; Miklos-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 Miklos-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 generalize the baseline model in two directions. First, we let demand 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 a form of countercyclicality similar to that in Rotemberg and Saloner.

Finally, we consider a generalized model that accommodates product differentiation and capacity constraints. The general 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 the details of the firms' payoff functions and can differ from that in the baseline, undifferentiated 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 available 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 the general form of an optimal disclosure policy is censorship of an intermediate interval of states, rather than upper censorship.

The remainder of the paper is organized as follows. Following a discussion of the literature, Section 2 presents the baseline model with iid demand or cost states. Section 3 solves the model and discusses its implication. Section 4 contains the extension to a persistent state. Section 5 contains the extension to a symmetric, linear demand system. Section 6 concludes. The general version of the model that allows a range of demand systems and capacity constraints is presented in Appendix A.

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

Much of the recent literature on pricing algorithms studies how independent algorithms can learn to set supra-competitive prices (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; Brown and MacKay, 2023; Hansen, Misra, and Pai, 2021; Lamba and Zhuk, 2024). This paper instead studies how a shared algorithm with demand information superior to the firms' optimally discloses information to facilitate collusion. Sugaya and Wolitzky (2018, Example 3) and Miklos-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. O'Connor and Wilson (2019), Martin and Rasch (2022), and Bonatti, Fiocco, and Piccolo (2024) document similar effects under imperfect monitoring.[3] 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 a firm that adopts and one that does not, rather than just the profit of adopters. Incorporating this consideration could be an interesting direction for future research. At the same time, Harrington (2024) also considers the case of coordinated adoption, which leads to a similar designer objective to ours.

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.[4] These papers find that

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

[4] The 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 papers on communicating private signals of a stochastic market demand state (Hanazono and Yang, 2007; Gerlach, 2009; Buehler and Gärtner, 2013; Sahueget and Walckiers, 2017).

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.[5]

Optimal information disclosure has been studied extensively in static environments (Rayo 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). Kolotilin and Li study a repeated cheap talk game with voluntary transfers. They show that the problem of characterizing the optimal equilibrium in this game can be reduced to a static information design problem, much as we do here.[6] 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.[7] 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 again to the special 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 4, and we also characterize more general optimal disclosure policies in Section 5 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 general idea is to consider static information design in the stage game augmented with continuation payoffs. This approach 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 undifferentiated goods or with differentiated goods with linear

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

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

[7]In 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.

demand), while the model in Appendix A is a relatively general one where this methodology applies.

2 A Model of Collusion with Information Disclosure

Prices and profits. We consider undifferentiated-product Bertrand competition among n firms with stochastic demand or a stochastic common cost of serving the market. In each period, a 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. A firm's information about s when setting its price 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 cost of serving the market, the corresponding assumptions are slightly different. Here, we assume that $\underline{s} \geq 0$ and $\Pi(p, s) \leq 0$ 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.[8]

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 $\mu \in \Delta(S)$, 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.[9]

[8] Another interpretation of the model is that the firms are bidders in a procurement auction, where the reserve price or the cost of fulfilling the contract is privately observed by an intermediary who coordinates bid-rigging among the firms. One example that fits this interpretation well is the Kumatori Contractors Cooperative studied by Kawai, Nakabayashi, and Ortner (2024).

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

Affineness in s is a strong assumption, but it is satisfied in some important cases. First, affineness holds if there is a binary underlying demand or cost state $\mathbf{s} \in \{\underline{s}, \bar{s}\}$ that is realized after prices are set, where s is a continuous signal of $\mathbf{s}$ satisfying $\Pr(\mathbf{s} = \bar{s}|s) = (s - \underline{s}) / (\bar{s} - \underline{s})$. Second, affineness holds in the canonical case of linear demand with an unknown intercept, where demand equals $D(p,s) = s - p$, and hence $\Pi(p,s) = p(s-p)$.[10] Third, affineness holds for linear demand with a known intercept normalized to 1 and an unknown per-unit production cost 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 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.[11] Since expected profit is measurable with respect to mean demand, it is without loss to view the intermediary as choosing a distribution G of the firms' posterior expectations of s. 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.[12]

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

[10] These 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 Miklos-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'.

[11] With 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 (2023) for results with concave payoffs, as well as a recent literature review.

[12] In Section 4, the state follows a Markov process, and the optimal disclosure policy depends on the previous period's state. In Section 5 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 $p_i(s)$ as a function of the realized disclosed mean demand state s and the history of past mean demand states and all firms' past prices.[13]

3 Optimal Information Disclosure

This section characterizes the disclosure policy and equilibrium that maximize collusive profits (the sum of the firms' payoffs). We reduce this problem to a static information design problem in Section 3.1 and then solve it in Section 3.2. We then discuss the model's empirical predictions, consumer welfare implications, and comparative statics in Section 3.3.

3.1 Reduction to Static Information Design

For any number $V \geq 0$, define

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

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

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

Note that the right-hand side of (2) is bounded and continuous in V by Berge's theorem, so V^* is well-defined by the intermediate value theorem.

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

[13] Restricting to pure strategies is standard but is not without loss of generality, as randomization can 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.

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) for $V = V^*$.*

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

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

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

Note that $p(s)$ is well-defined by the intermediate value theorem, as $\Pi(0, s) = 0$ and $\Pi(p, s)$ is continuous in p.[14] Let $G^* \in \text{argmax}_{G \in MPC(F)} \mathbb{E}^G[\min\{\Pi^m(s), \Pi^{\max}(\delta, n, V^*)\}]$. Consider disclosure policy G^*, together with the strategy profile where all firms price at $p(s)$ whenever mean demand s 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 s is to price just below $p(s)$: this is immediate if $p(s) = p^m(s)$ and otherwise follows because $p(s)$ is the smallest price p satisfying $\Pi(p, s) = \Pi^{\max}(\delta, n, V^*)$, so that $\Pi(p', s) < \Pi^{\max}(\delta, n, V^*)$ for all $p' < p(s)$. 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 an equilibrium if and only if, for all s, we have

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

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

[14] Here 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.

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 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 s, the winning price is $p(s)$ and each firm i wins with probability α_i and obtains an equilibrium continuation value of v_i (we use lower-case letters for per-firm values). (So, $\alpha_i = 1/\left|j : p_j(s) = p(s)\right|$ if $p_i(s) = p(s)$, and $\alpha_i = 0$ otherwise. Note that each $p_i(s)$—and thus the winning price $p(s)$—is deterministic by our restriction to pure strategy equilibria.) Since a possible deviation by firm i is to price just below $p(s)$, firm i's incentive constraint implies

$$(1-\delta)\Pi(p(s), s) \le \alpha_i (1-\delta)\Pi(p(s), s) + \delta v_i.$$

Averaging this inequality over the n firms, we have

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

where the second inequality is by definition of $\bar{V}$. Therefore, $\Pi(p(s), s) \le \Pi^{\max}\left(\delta, n, \bar{V}\right)$, and hence expected collusive profits in period t are at most $\max_{G \in MPC(F)} \mathbb{E}^G\left[\min\left\{\Pi^m(s), \Pi^{\max}\left(\delta, n, \bar{V}\right)\right\}\right]$. Since this holds for every period t, we have $\bar{V} \le \max_{G \in MPC(F)} \mathbb{E}^G\left[\min\left\{\Pi^m(s), \Pi^{\max}\left(\delta, n, \bar{V}\right)\right\}\right]$. But this implies that $\bar{V} \le V^*$, 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 (1986). Conversely, if $\delta \ge (n-1)/n$ then monopoly profit under no information disclosure, $\Pi^m\left(\mathbb{E}^F[s]\right)$, is attainable.

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

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 $\delta \ge (n-1)/n$ then no information disclosure together with a constant on-path price of $p^m\left(\mathbb{E}^F[s]\right)$ and zero prices off path is an equilibrium, with expected profit $\Pi^m\left(\mathbb{E}^F[s]\right)$. ■

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

3.2 Solving the Information Design Problem

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

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

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

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 $V^* \geq \Pi^m\left(\mathbb{E}^F[s]\right)$, and hence $s^* \geq \mathbb{E}^F[s]$.[15] Thus, there exists $\hat{s} \in [\underline{s}, \bar{s}]$ such that

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

We can now characterize the optimal disclosure policy and optimal collusive prices. Figure 1 illustrates the optimal 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 unique optimal collusive price $p(s)$ in state s is given by*

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

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 $\hat{s}$.*[16] *The unique optimal collusive price $p(s)$ in state s is given by (4) with the reversed inequalities.*

---

[15] This follows because if $s^* < \mathbb{E}^F[s]$ then $\Pi^m\left(\mathbb{E}^F[s]\right) > \Pi^{\max}(\delta, n, V^*)$ 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^*.

[16] Here, the definitions of $\Pi^{\max}$ and s^* remain as in (1) and (3).

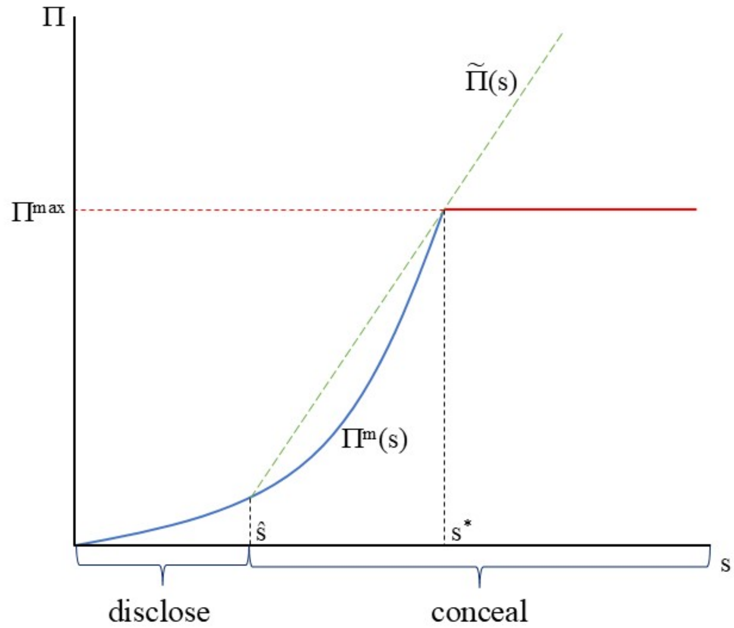

Figure 1: The optimal disclosure policy. First, s^* is determined as the solution to $\Pi^m(s^*) = \Pi^{\max}(\delta, n, V^*)$. Then, $\hat{s}$ is determined as the solution to $E^F[s|s \geq \hat{s}] = s^*$. The optimal information policy discloses demand states $s < \hat{s}$ and recommends the corresponding 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 auxiliary objective function $\widetilde{\Pi}(s)$ is defined in the proof of Theorem 1.

With stochastic demand, note that $\hat{s} = \underline{s}$—so it is optimal to reveal nothing about demand—iff $s^* = \mathbb{E}^F[s]$, which holds iff $\delta = (n-1)/n$. Conversely, $\hat{s} = \bar{s}$—so it optimal to fully disclose demand—iff $s^* = \bar{s}$, which holds iff $\Pi^m(\bar{s}) \geq \Pi^{\max}(\delta, n, V^*)$. Otherwise, we have $\Pi^m\left(\mathbb{E}^F[s]\right) < \Pi^{\max}(\delta, n, V^*) < \Pi^m(\bar{s})$, so partial disclosure is optimal.

The intuition for Theorem 1 is that disclosing demand information increases expected monopoly profits—as $\Pi^m(s)$ is convex—but revealing that the expected state is too high requires cutting price to deter a deviation (as in Rotemberg and Saloner, 1986). 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.

Proof. Recall that $\Pi^m(s)$ is convex in s, while $\Pi^{\max}(\delta, n, V^*)$ is independent of s. Thus, $\min\{\Pi^m(s), \Pi^{\max}(\delta, n, V^*)\}$ is increasing and convex in s for $s \leq s^*$ and is constant in s for $s > s^*$. In particular, $\min\{\Pi^m(s), \Pi^{\max}(\delta, n, V^*)\}$ is "S-shaped" in s: first convex, then concave. So, (2) describes a mean-measurable information design problem with an S-shaped objective function. 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 just 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.

We sketch the proof for completeness.[17] 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.) 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

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

Since $\tilde{\Pi}(s) \geq \min\{\Pi^m(s), \Pi^{\max}(\delta, n, V^*)\}$ for all s, this is an upper bound for $\max_{G \in MPC(F)} \mathbb{E}^G\left[\min\{\Pi^m(s), \Pi^{\max}(\delta, n, V^*)\}\right]$. 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. ■

3.3 Implications and Comparative Statics

We now discuss the implications of Theorem 1 for firm profits, prices, and consumer surplus. In discussing the price implications, we make the standard assumption that the monopoly price $p^m(s)$ is unique and increasing in s. For example, this holds if $\Pi(p, s)$ is strictly quasi-concave in p and supermodular in (p, s).

[17] Our proof follows 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 B.

Price rigidity—not price wars—during booms. We first compare collusive prices under optimal disclosure and full disclosure (i.e., in Rotemberg and Saloner). Optimal disclosure leads to price rigidity during booms, rather than price wars as in Rotemberg and Saloner. To see this, let V^{FD} be optimal collusive profit under full disclosure, which is given by the greatest fixed point of the equation

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

and let s^{FD} solve

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

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

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

Since $p^m(s)$ is increasing and $\Pi(p, s)$ is continuous in p and 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.

In contrast, for the optimal collusive prices under the optimal disclosure policy, (4), we see that $p(s)$ is increasing for $s < s^*$ and constant (at a higher level) for $s \geq s^*$. Thus, optimal disclosure entails a policy of price rigidity at demand states $s \geq s^*$, rather than "price wars" as in Rotemberg and Saloner's model. This result gives a novel rationale for oligopoly price rigidity: prices are rigid because colluding firms optimally limit their own information about market demand to deter deviations.[18]

Supra-monopoly pricing. A notable implication of Theorem 1 is that optimal collusive prices are *above monopoly* at intermediate demand states: for $s \in (\hat{s}, s^*)$, the op-

[18] Carlton (1986) and others find that prices are more rigid in concentrated industries, and Harrington (2008) and others suggest price rigidity as a possible 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 privately informed firms to reveal their 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).

timal 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 could be attained at any one of these states by disclosing the state s and recommending price $p^m(s)$ (holding the rest of the equilibrium fixed). That is, 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 all demand states above $\hat{s}$ results in an inefficiently high price for demand states in $s \in (\hat{s}, s^*)$, but thereby supports a higher price for demand states $s > s^*$ than would be attainable under full information.[19]

Higher prices in all demand states; lower consumer surplus. We next show that prices are higher for every demand state s under optimal disclosure as compared to full disclosure.

Proposition 1 *As compared to collusive prices under full disclosure, collusive prices under the optimal disclosure policy are higher at each demand state.*

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 $p(s) = p^m(s^*) \geq p^m\left(s^{FD}\right) \geq p^{FD}(s)$, where the first inequality follows because $s^{FD} \leq s^*$ and p^m is increasing, and the second follows because $s^{FD} \leq s$ and p^{FD} is decreasing to the right of s^{FD}. ■

In particular, Proposition 1 implies that consumer surplus is lower under optimal disclosure as compared to full disclosure. That is, the firms' ability to limit their own information via an algorithm unambiguously harms consumers.

Comparative statics. We now turn to comparative statics, starting with comparative statics for the number of firms n and the discount factor δ. In what follows, we say that

[19] Supra-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).

prices are *more flexible* if $\hat{s}$ is higher, so that a wider range of demand states are disclosed and $p(s) = p^m(s)$ for a wider range of states.[20]

Proposition 2 *A decrease in the number of firms n or an increase in the discount factor δ increases collusive profit V^* and makes collusive prices more flexible.*

Proof. Note that n and δ affect V^* and p only through the function $\Pi^{\max}$, which is decreasing in n and increasing in δ. Thus, decreasing n or increasing δ shifts the right-hand side of (2) up 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. ■

Among the results in Proposition 2, the prediction that increasing n makes prices more rigid is the most distinctive to our model. The logic is that increasing n makes incentive constraints bind for a wider range of demand state realizations, which necessitates pooling a wider range of states to deter deviations.

Another interesting comparative statics question concerns the impact of improving the algorithm's information (or "accuracy"). By Blackwell (1953), this corresponds to taking a mean-preserving spread of F, which expands the set $MPC(F)$ of implementable distributions of posterior mean demands. It is immediate that improving the algorithm's accuracy increases optimal collusive profit.

A more subtle question concerns the effect of improving the algorithm's accuracy on consumer surplus. To address this, let $CS(s)$ denote consumer surplus in demand state s at the monopoly price $p^m(s)$. Under full disclosure and monopoly pricing (which is implementable if δ is high enough), improving the algorithm's accuracy increases expected consumer surplus if $CS(s)$ is convex and decreases expected consumer surplus if $CS(s)$ is concave. The latter case is the standard one: for example, $CS(s)$ is concave under linear demand with either an unknown demand intercept (i.e., $\Pi(p,s) = p(s-p)$) or an unknown cost (i.e., $\Pi(p,s) = (p-s)(1-p)$).[21] We now show that in this case, improving the algorithm's accuracy also decreases expected consumer surplus under the optimal information disclosure

[20] Within the class of optimal censorship disclosure policies, increasing $\hat{s}$ is equivalent to making the algorithm more informative in the Blackwell order.

[21] This holds because in the former case $p^m(s) = s/2$ and hence $CS(s) = s^2/8$, and in the latter case $p^m(s) = (1+s)/2$ and hence $CS(s) = (1-s)^2/8$.

policy. The logic is that a more accurate algorithm optimally discloses more information to firms while recommending the corresponding monopoly price at each disclosed mean demand state, so the effect of improving the algorithm's accuracy is the same as the effect of improving a monopoly firm's information.

Proposition 3 *Assume that consumer surplus under monopoly pricing, $CS(s)$, is concave in s. Then improving the algorithm's accuracy reduces expected consumer surplus.*

Proof. By Theorem 1, the collusive equilibrium price at any disclosed mean demand state s is the corresponding monopoly price $p^m(s)$. To prove the proposition, it thus suffices to show that for any distributions (F_1, F_2, G_1, G_2) where F_2 is a mean-preserving spread of F_1, G_1 is the distribution of s under an optimal disclosure policy for prior F_1, and G_2 is the distribution of s under an optimal disclosure policy for prior F_2, we have that G_2 is a mean-preserving spread of G_1. We defer the proof of this fact to the appendix. ■

Proposition 3 shows that improving the algorithm's accuracy reduces consumer surplus whenever consumer surplus is concave in demand states under monopoly pricing. This finding contrasts with results of Sugaya and Wolitzky (2018, Example 3) and Miklos-Thal and Tucker (2019), who find that a more accurate demand prediction algorithm reduces consumer surplus when the firms' discount factor δ lies in an intermediate range. The reason for this difference is that these papers assume that the algorithm fully discloses its information to the firms, which tightens incentive constraints and thus necessitates a reduction in equilibrium prices when δ is not too high. In contrast, with optimal information disclosure, a more accurate algorithm always increases average prices, and Proposition 3 shows that it also reduces consumer surplus whenever consumer surplus is concave under monopoly pricing. Our assessment of the likely impact of improved algorithmic demand prediction on consumer surplus is thus considerably more pessimistic than that in prior work.

Empirical implications, collusive markers, and the interpretation of "price wars." We close this section with a brief discussion of the model's empirical implications. There are four 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 price and monopoly price, $p^m(s) - p(s)$, is non-monotone: first zero, then negative (when price is supra-monopolistic), then positive.

4. Prices are more flexible (and higher on average) when firms are more patient or fewer in number.

We offer two remarks on these empirical implications.

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

Second, 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: price wars occur in low demand states as part of an optimal repeated game equilibrium under imperfect monitoring of competitors' prices. In contrast, in Rotemberg and Saloner (1986), prices are counter-cyclical in higher demand states: price wars occur in high demand states due to perfect monitoring and demand information. Interestingly, while our model is much closer to Rotemberg and Saloner's than to Green and Porter'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 selectively disclosed demand information 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 price and monopoly price, $p^m(s) - p(s)$, over the cycle. In Green and Porter's theory, the gap is larger in low demand states: collusion is "more successful" when demand is high. In our theory, 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 these competing predictions.

Some narrative evidence on this comparison comes from the RealPage case discussed in the Introduction. Consistent with our prediction of more successful collusion in low demand states, RealPage seems to have taken particular pride in its performance in down markets, especially during the Covid pandemic. For example, the US DOJ complaint against RealPage asserts that, "in down markets… [RealPage] instills pricing discipline in landlords, curbing normal fully independent competitive reactions by substituting them with interdependent decision-making," (USDOJ 2024a, p. 48). In its own words, RealPage advertised that the "AI and the robust data in the RealPage ecosystem" helps its clients "avoid the race to the bottom in down markets," (ibid., p. 46). These statements seem consistent with our prediction of flexible monopoly prices $p^m(s)$ in low demand states and a rigid, non-monopoly price $p^m(s^*)$ in "normal times." Another example comes from the bidding ring organized by the Kumatori Contractors Cooperative, studied by Kawai, Nakabayashi, and Ortner (2024). Notably, this organization took drastic steps to limit bidders' information about the cost of completing the largest construction project they bid on (ibid., pp. 24–25), which matches our result that the intermediary censors information in high-demand states where deviating is most tempting.

4 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

[22] With 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 very different predictions about prices, $p(s)$ (non-decreasing in our model; single-peaked in theirs), but more 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).

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 profit is measurable with respect to mean demand (or cost, but we continue to focus on the demand case). To do so, we require 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 assumption is realistic if firms observe their sales at the end of the 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:

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

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 is F_{s_0} for some $s_0 \in [\underline{s}, \bar{s}]$.

Note that 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 of these cases are of interest: positive persistence is arguably more natural, but negative persistence has been used to capture cyclical demand movements, for example by Haltiwanger and Harrington (1991).

The characterization of the optimal disclosure policy and optimal collusive prices are the same as in the iid case, except that now the expected present value of collusive profits $V(s)$ 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:

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

Note that 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

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

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 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 (2) for $V(\cdot) = V^*(\cdot)$.*

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.[23]

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), $\Pi^m\left(\min\left\{\mathbb{E}^{F_{\underline{s}}}[s], \mathbb{E}^{F_{\bar{s}}}[s]\right\}\right)$, is attainable for any initial state.

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

[23] The 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.

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 $\min\left\{p : \Pi\left(p, \mathbb{E}^{F_s}[\tilde{s}]\right) = \Pi^m\left(\min\left\{\mathbb{E}^{F_{\underline{s}}}[\tilde{s}], \mathbb{E}^{F_{\bar{s}}}[\tilde{s}]\right\}\right)\right\}$ when the previous period demand state is s (noting that this price is well-defined by the intermediate value theorem, as $\Pi\left(p, \mathbb{E}^{F_s}[\tilde{s}]\right)$ is continuous in p and monotone in 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.[24] 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 $\hat{s}(s) = \underline{s}$ otherwise. Note that, by Lemma 4 and our assumption that $\delta \geq (n-1)/n$, we have $V^*(\underline{s}) \geq \Pi^m\left(\min\left\{\mathbb{E}^{F_{\underline{s}}}[s], \mathbb{E}^{F_{\bar{s}}}[s]\right\}\right)$, and hence $s^* \geq \min\left\{\mathbb{E}^{F_{\underline{s}}}[s], \mathbb{E}^{F_{\bar{s}}}[s]\right\}$, 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 $\min\left\{p : \Pi\left(p, \mathbb{E}^{F_s}[s']\right) = \Pi^{\max}\left(\delta, n, W^*\left(\mathbb{E}^{F_s}[s']\right)\right)\right\}$, which is less than the corresponding monopoly price $p^m\left(\mathbb{E}^{F_s}[s']\right)$.

We are now prepared to characterize the optimal disclosure policy and optimal collusive prices 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*

[24] There 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 $\Pi^m(\underline{s}) \leq \Pi^{\max}(\delta, n, \Pi^m(\underline{s})) < \Pi^{\max}\left(\delta, n, \Pi^m(\min\left\{\mathbb{E}^{F_{\underline{s}}}[s], \mathbb{E}^{F_{\bar{s}}}[s]\right\})\right) \leq \Pi^{\max}(\delta, n, W^*(\underline{s}))$, and the left-hand side of (7) is convex while the right-hand side is linear.

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*

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

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* $\hat{s}(s) = \bar{s}$ *in the third line. 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 satisfying $\Pi\left(p, \mathbb{E}^{F_s}[s']\right) = \Pi^{\max}\left(\delta, n, W^*\left(\mathbb{E}^{F_s}[s']\right)\right)$. 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. The key point is that, whenever the optimal censorship policy is non-trivial (i.e., $\hat{s}(s) \in (\underline{s}, \bar{s})$, so that some states are disclosed and others are concealed), the mean demand among the concealed states must equal s^*, regardless of the last-period demand state s. Thus, with positive persistence, $\hat{s}(s)$ is decreasing, as F_s is higher when s is higher, so more low states must be pooled in with higher states to induce the constant mean demand s^* among concealed states. So, with positive persistence, the algorithm discloses less information in good times, when firms are optimistic about demand and are thus more inclined to deviate. Conversely, with negative persistence, the algorithm discloses more information in good times, when firms are pessimistic and thus less inclined to deviate.[25]

We also note that, in contrast to the iid case, it can now be generically optimal to reveal nothing about the current-period demand state.[26] 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 latter case where $\hat{s}(s) = \underline{s}$, the optimal price $p\left(\mathbb{E}^{F_s}[s'];s\right)$ satisfies $\Pi\left(p, \mathbb{E}^{F_s}[s']\right) = \Pi^{\max}\left(\delta, n, W^*\left(\mathbb{E}^{F_s}[s']\right)\right)$, which is less than the corresponding monopoly price $p^m\left(\mathbb{E}^{F_s}[s']\right)$ and can even be decreasing in s over some range. Thus, while optimal prices are always monotone in current demand (as in the iid case and in contrast to Rotemberg Saloner), they are not necessarily monotone in the 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.[27]

In addition to these more novel points, Theorem 2 shows that the results of the iid model generalize to the Markov case. For example, for any last-period demand state s, optimal collusion entails price rigidity at high current demand states, supra-monopoly prices over an intermediate range of states, and more rigid prices when δ is lower or n is higher. Proposition 3 also extends to the Markov case, noting that improving the algorithm's accuracy now corresponds to taking a mean-preserving spread of each distribution F_s.

A final question concerns the comparative statics of optimal disclosure, firm profit, and

[25] The results for the stochastic cost case are analogous. Now, with positive persistence, the algorithm discloses less information when the last-period cost was low; and with negative persistence, the algorithm discloses more information when the last-period cost was low.

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

[27] Whether 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 $\{s : \hat{s}(s) = \underline{s}\}$.

consumer surplus as demand becomes more persistent. This is a complex question, so we just provide a numerical example that illustrates some interesting possibilities in Appendix C. In particular, the example shows that the effect of greater persistence of demand on the amount of information disclosure, collusive profit, and consumer surplus can all be non-monotone.

5 Differentiated Products

We now extend the baseline iid model from the undifferentiated goods setting of Rotemberg and Saloner (1986) 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, 2024) 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 also include as a special case the baseline, undifferentiated goods model, with or without capacity constraints.

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

$$\pi_i(\mathbf{p}, s) = p_i \left(s + c \sum_{j \neq i} p_j - p_i \right) - p_i^2$$

for the stochastic demand case, or

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

for the stochastic cost case.[28] Note that $\pi_i(\mathbf{p}, s)$ is affine in s, so our key assumption is satisfied. For concreteness, we focus on the stochastic demand case. (In general, we use Π to denote the industry profit and π to denote an individual firm's profit.)

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

[28] The condition $c \geq 0$ says that the goods are substitutes. The condition $c \leq 1/(n-1)$ 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).

at public mean belief s as

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

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

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

Notice that $\underline{\pi}(s)$ 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 demand. Thus, the *no-information competitive outcome*, where the algorithm discloses no information and expect profit equals

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

minimizes the firms' static Nash profit.

We will 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 undifferentiated goods case where the worst static Nash payoff equals the minimax payoff of 0, this off-path payoff $\underline{\pi}$ is strictly greater than the minimax payoff of $\mathbb{E}^F[s]^2/4$. In other words, we focus on equilibria sustained by the threat of *Nash reversion*, which entails a loss of optimality in the class of all pure-strategy, subgame perfect equilibria. To find the optimal equilibrium in this class, one would have to 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 exactly the same characterization of optimal equilibrium prices and information disclosure. Our qualitative results are thus insensitive to the specification of off-path payoffs.

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

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

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

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

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

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

which is given by

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

and let

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

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 payoff $\underline{\pi}$, and $\pi^{\max}(s,v)$ is the corresponding per-firm profit in state s. (Again, a small letter v is the per-firm value, while V is the industry value.) 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 c). This contrasts with the undifferentiated goods case, where $p^{\max}(s,v)$ is decreasing and $\pi^{\max}(s,v)$ is constant in s for $s>s^{*}$.

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

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

Lemma 1 extends to the current setting: in particular, optimal collusive profit equals v^{*}. Formally, this follows from the more general Lemma 5 in Appendix A. The key step in the argument is that symmetric pricing remains optimal in each demand state, which holds because, for any price vector $\mathbf{p} = (p_1, \ldots, p_n)$, we have

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

(as can be verified by straightforward calculation, using the assuming that $c \geq 0$), so replacing any asymmetric vector $\mathbf{p}$ with the constant price vector $(\sum_i p_i/n, \ldots, \sum_i p_i/n)$ 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

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

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

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

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 is true because, since $\pi^m(s)$ and $\pi^{\max}(s, v^*)$ are convex in s, if $(\hat{s}_L, \hat{s}_H)$ 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 $\frac{s_H - s^*}{s_H - s_L}\pi^m(s_L) + \frac{s^* - s_L}{s_H - s_L}\pi^{\max}(s_H, v^*) < \pi^m(s^*)$, and 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 $\frac{s_H - s^*}{s_H - s_L}\pi^m(s_L) +$

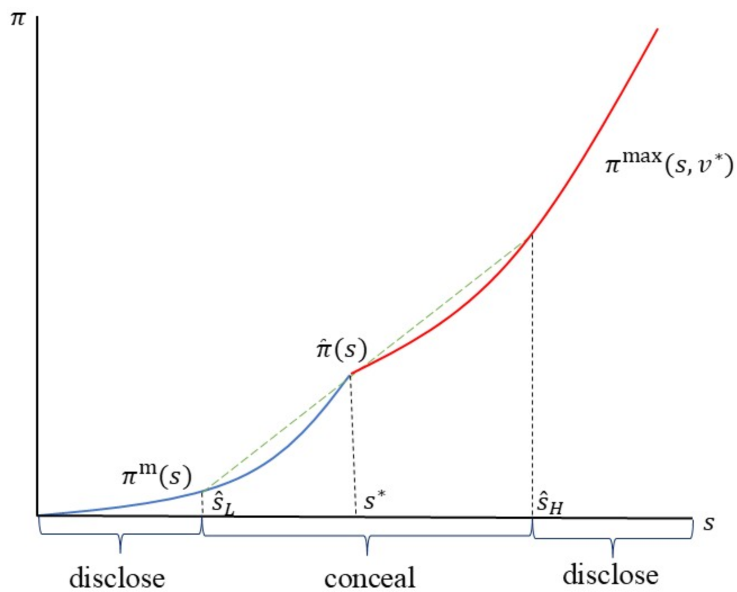

Figure 2: The optimal disclosure policy with a symmetric, linear demand system. The cutoffs $(\hat{s}_L, \hat{s}_H)$ are defined in the text. The optimal information policy discloses demand states $s < \hat{s}_L$ and recommends the corresponding monopoly price, $p^m(s)$; conceals demand states $s \in [\hat{s}_L, \hat{s}_H]$ 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.

$\frac{s^*-s_L}{s_H-s_L}\pi^{\max}(s_H, v^*) > \pi^m(s^*)$.

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. The economic 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 on on-path equilibrium profit (as in the undifferentiated goods case), and indeed 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 undifferentiated 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 necessarily the highest demand states that are censored: depending on parameters, it could instead be optimal to censor an intermediate interval of states, or only the lowest states.

6 Conclusion

This paper has developed a simple 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 Kalibrate in retail gasoline, 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. Under the assumption that expected profits are determined by the expected demand state, we show that 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 considerable *price rigidity*, as well as *supra-monopoly prices* for a range of intermediate demand states. We also establish several comparative statics results: prices are more flexible (and higher on average) when firms are more patient or fewer in number; a more accurate algorithm reduces expected consumer surplus under a natural concavity condition; and more demand information is disclosed when recent demand was lower if demand is positively serially correlated. The second of these results gives reason for pessimism regarding the likely effect of improved algorithmic demand prediction on consumer surplus, in contrast to prior studies that find more optimistic results under the hypothesis that the algorithm always discloses its predictions (Sugaya and Wolitzky, 2018; Miklos-Thal and Tucker, 2019). Finally, we extend the model to allow product differentiation or capacity constraints: here, most of the analysis goes through, although the specific form of the optimal censorship region depends on the demand system.

Many of our assumptions could be further relaxed at the cost of a more intricate analysis. First, if expected profit depends on the entire distribution of the unknown demand 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 will be approximately optimal if the information design problem is close to linear. Second, if the intermediary's objective differs from maximizing industry profit, as in Harrington (2022), or if some firms do not use the intermediary, the model must be extended to incorporate for the intermediary's incentives and firms' incentives to use the intermediary. Finally, allowing asymmetric, private information disclosure by the algorithm (or exogenous asymmetric, private information for the firms) appears challenging but potentially insightful. 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 5, let $\pi_i(\mathbf{p}, s)$ denote firm i's profit at price vector $\mathbf{p} = (p_1, \ldots, p_n)$ and demand state s.[29] We also define firm i's maximum deviation profit at price vector $\mathbf{p}$ by

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

We impose four key assumptions on payoffs. First, we assume that payoffs are symmetric, continuous, and quasi-concave in a firm's own price.[30]

Assumption 1 For any price vector $(p_1, \ldots, p_n)$, state s, firm i, and permutation ϕ on $\{1, \ldots, n\}$, we have $\pi_i\left(p_{\phi(1)}, \ldots, p_{\phi(n)}, s\right) = \pi_{\phi(i)}(p_1, \ldots, p_n, s)$. 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 equilibrium at mean public belief $\mathbb{E}^F[s]$. As in Section 5, for concreteness we will 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 $\pi(p, s) = \pi_i(p, \ldots, p, s)$, and we define the corresponding maximum deviation profit as $\pi^d(p, s) = \pi_i^d(p, \ldots, p, s)$. 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.

[29] This appendix focuses on stochastic demand. The stochastic cost case is analogous.

[30] The 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 $\{1, \ldots, n\}$ such that $\pi(i) = j$ and, for any price vector $(p_1, \ldots, p_n)$, state s, and firm k, we have $\pi_i\left(p_{\phi(1)}, \ldots, p_{\phi(n)}, s\right) = \pi_{\phi(k)}(p_1, \ldots, p_n, s)$. The examples in the current paper are all totally symmetric, but many oligopoly models (e.g., Salop's circle model) are only weakly symmetric. See Plan (2023).

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

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

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 5. 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.[31]

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

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

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

[31] If s is continuous and $D(p,s) = s - p$, then $\pi_i(\mathbf{p},s)$ is linear up to a choke price. As in footnote 10, 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'.

point of the equation

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

which is now well-defined by Tarski's theorem.

Lemma 1 extends as follows.

Lemma 5 *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) for $v = v^*$*

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

We now show that no equilibrium can attain higher profits. 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

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

Averaging this inequality over the n firms, we have

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

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

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

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

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

Moreover, since $p = p^{\max}(s, \bar{v})$ maximizes $\pi(p, s)$ subject to $\pi^d(p, s) - \pi(p, s) \leq (\delta/(1-\delta))\bar{v}$ (by Assumption 4), $p^{\max}(s, \bar{v})$ satisfies these inequalities. Hence, replacing $\mathbf{p}(s)$ with the constant price vector $(p^{\max}(s, \bar{v}), \ldots, p^{\max}(s, \bar{v}))$ for each state s in period t increases profits in period t without violating incentive constraints in period t or any earlier period. Therefore, expected collusive profits in period t are 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 results 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.[32]

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 5. (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 Assumption 1–3 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*

[32] It 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

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

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 binds off path but is slack on-path: C is greater than the per-firm monopoly quantity at the highest demand state, $\bar{s}/(2n)$, and smaller than the industry quantity at the lowest demand state $\underline{s}/2$ with the monopoly price. 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

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

This expression is convex in s, so Condition 1 holds, and, in general, a disclose-conceal-disclose disclosure policy is optimal, similarly to the differentiated goods setting in Section 5. 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.

B Proof of Proposition 3

Since F_2 is a mean-preserving spread of F_1, we have $\mathbb{E}^{F_1}[s] = \mathbb{E}^{F_2}[s]$ and $\int_{\underline{s}}^{s} F_1(s)\,ds \leq \int_{\underline{s}}^{s} F_2(s)\,ds$ for all s. By Theorem 1, we have

$$G_1(s) = \begin{cases} F_1(s) & \text{if } s \leq \hat{s}_1, \\ F_1(\hat{s}_1) & \text{if } \hat{s}_1 < s < s_1^*, \\ 1 & \text{if } s \geq s_1^*, \end{cases}$$

where $s_1^* = E^{F_1}[s|s > \hat{s}_1]$, and similarly for G_2. In addition, collusive profit is higher under G_2 than G_1 (as improving the algorithm's information increases collusive profit), which implies that $s_1^* \leq s_2^*$, by (3).

We show that these conditions imply that $\int_{\underline{s}}^{s} G_1(s)\,ds \leq \int_{\underline{s}}^{s} G_2(s)\,ds$ for all s, showing that G_2 is a mean-preserving spread of G_1. It suffices to consider the case where there exists

s^* such that $s_1^* = s_2^* = s^*$, as increasing s_2^* corresponds to taking another mean-preserving spread of G_2. Note that, for any $s \leq \hat{s}_2$, we have

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

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 $s < s_1^*$, and the second inequality is because F_2 is a mean-preserving spread of F_1. 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, $\int_{\underline{s}}^{s} (G_1(s) - G_2(s))\, ds$ 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 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

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

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.

C Example with Persistent Demand

We assume a binary demand state—so $s \in \{\underline{s}, \bar{s}\}$—and linear demand—so $\Pi(p, s) = p(s - p)$—and we 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. Specifically, 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 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 the continuation value at last-period state $\bar{s}$ relaxes the binding incentive constraint. In contrast, the effect of increasing ρ on profits in 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

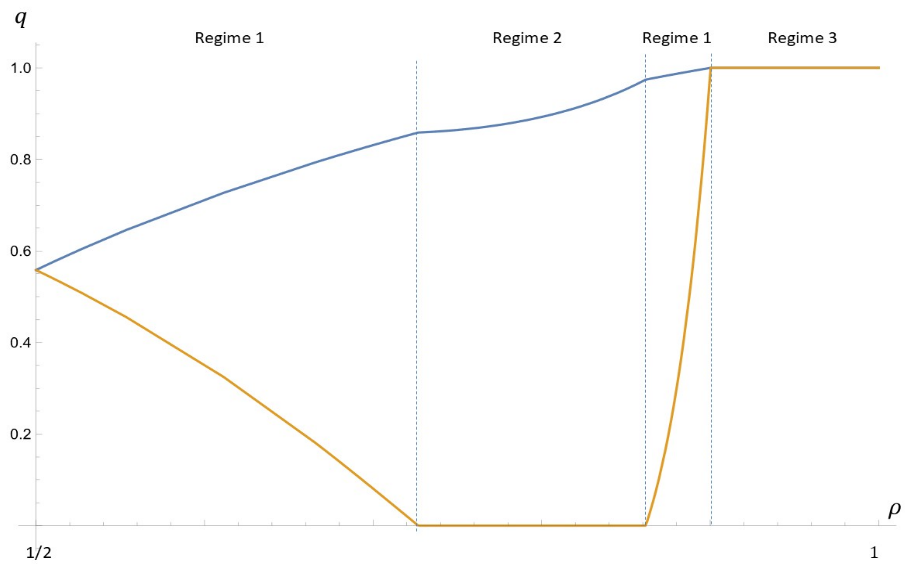

Figure 3: Optimal disclosure policy. The blue curve is the probability of disclosing $\underline{s}$ when the current state is $\underline{s}$ and the last-period state was $\underline{s}$; the orange curve is the probability of disclosing $\underline{s}$ when the current state is $\underline{s}$ and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

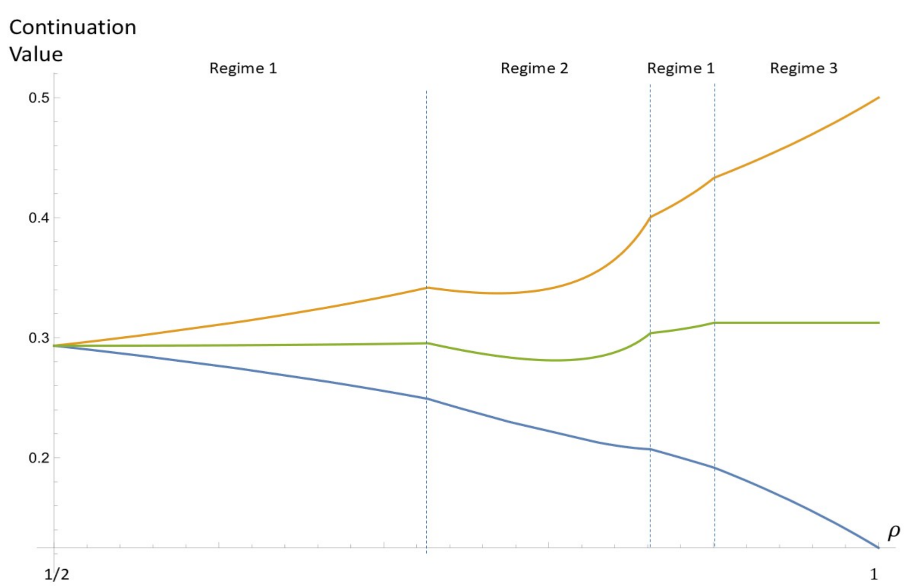

Figure 4: Optimal continuation values. The blue curve is a firm's continuation value at last-period state $\underline{s}$. The orange curve is the value at last-period state $\bar{s}$. The green curve is the ex ante expected profit.

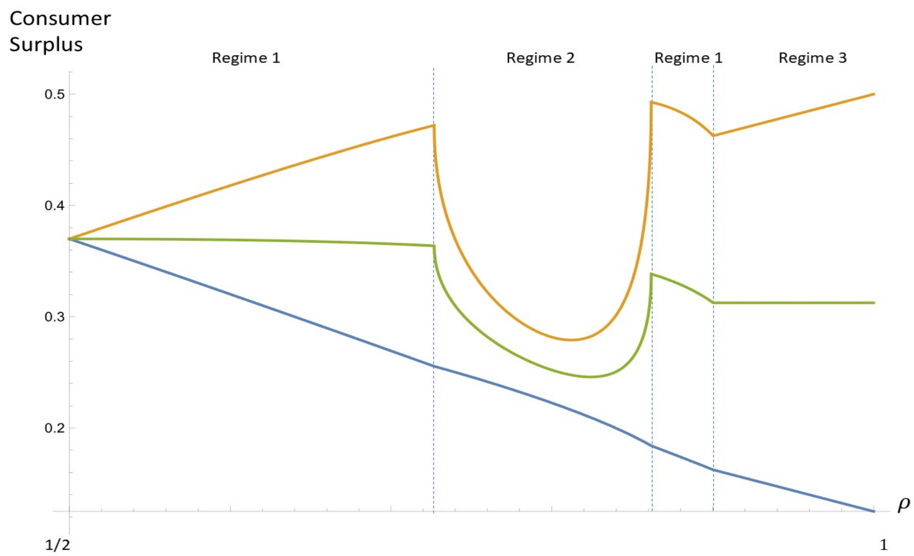

Figure 5: Consumer surplus. The blue curve is the current-period consumer surplus at last-period state $\underline{s}$. The orange line is the corresponding value at last-period state $\bar{s}$. The green curve is ex ante expected consumer surplus.

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] Brown, Zach Y., and Alexander MacKay. "Competition in Pricing Algorithms." *American Economic Journal: Microeconomics* 15.2 (2023): 109-156.

[16] Buehler, Stefan, and Dennis L. Gärtner. "Making Sense of Nonbinding Retail-Price Recommendations." *American Economic Review* 103.1 (2013): 335-359.

[17] 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).

[18] Calvano, Emilio, et al. "Protecting Consumers from Collusive Prices due to AI." *Science* 370.6520 (2020a): 1040-1042.

[19] Calvano, Emilio, et al. "Artificial Intelligence, Algorithmic Pricing, and Collusion." *American Economic Review* 110.10 (2020b): 3267-3297.

[20] Carlton, Dennis. "The Rigidity of Prices." *American Economic Review* 76.4 (1986): 637-58.

[21] 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.

[22] Dworczak, Piotr, and Giorgio Martini. "The Simple Economics of Optimal Persuasion." *Journal of Political Economy* 127.5 (2019): 1993-2048.

[23] Ellison, Glenn. "Theories of Cartel Stability and the Joint Executive Committee." *Rand Journal of Economics* 25.1 (1994): 37-57.

[24] Ely, Jeffrey C. "Beeps." *American Economic Review* 107.1 (2017): 31-53.

[25] Ezrachi, Ariel, and Maurice E. Stucke. "Artificial Intelligence & Collusion: When Computers Inhibit Competition." *University of Illinois Law Review* (2017): 1775.

[26] Gentzkow, Matthew, and Emir Kamenica. "A Rothschild-Stiglitz Approach to Bayesian Persuasion." *American Economic Review* 106.5 (2016): 597-601.

[27] Gerlach, Heiko. "Stochastic Market Sharing, Partial Communication and Collusion." *International Journal of Industrial Organization* 27.6 (2009): 655-666.

[28] Green, Edward J., and Robert H. Porter. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52.1 (1984): 87-100.

[29] 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.

[30] 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.

[31] Hanazono, Makoto, and Huanxing Yang. "Collusion, Fluctuating Demand, and Price Rigidity." *International Economic Review* 48.2 (2007): 483-515.

[32] Harrington Jr, Joseph E. "How do Cartels Operate?." *Foundations and Trends in Microeconomics* 2.1 (2006): 1-105.

[33] Harrington Jr, Joseph E. "Detecting Cartels." *Handbook of Antritrust Economics*, Buccirossi ed. (2008).

[34] Harrington Jr, Joseph E. "The Effect of Outsourcing Pricing Algorithms on Market Competition." *Management Science* 68.9 (2022): 6889-6906.

[35] Harrington Jr, Joseph E. "An Economic Test for an Unlawful Agreement to Adopt a Third-Party's Pricing Algorithm." *Economic Policy* (2024): eiae054.

[36] Harrington, Joseph E., and Andrzej Skrzypacz. "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices." *American Economic Review* 101.6 (2011): 2425-2449.

[37] Hickok, Nathaniel. "Algorithm Design Meets Information Design: Price Recommendation Algorithms on Online Platforms." *Working Paper*, MIT.

[38] Kamenica, Emir, and Matthew Gentzkow. "Bayesian Persuasion." *American Economic Review* 101.6 (2011): 2590-2615.

[39] Kandori, Michihiro. "The Use of Information in Repeated Games with Imperfect Monitoring." *Review of Economic Studies* 59.3 (1992): 581-593.

[40] Kandori, Michihiro. "Correlated Demand Shocks and Price Wars during Booms." *Review of Economic Studies* 58.1 (1991): 171-180.

[41] 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* (2024): Forthcoming.

[42] Kolotilin, Anton. "Optimal Information Disclosure: A Linear Programming Approach." *Theoretical Economics* 13.2 (2018): 607-635.

[43] Kolotilin, Anton, et al. "Persuasion of a Privately Informed Receiver." *Econometrica* 85.6 (2017): 1949-1964.

[44] Kolotilin, Anton, and Hongyi Li. "Relational Communication." *Theoretical Economics* 16.4 (2021): 1391-1430.

[45] Kolotilin, Anton, Timofiy Mylovanov, and Andriy Zapechelnyuk. "Censorship as Optimal Persuasion." *Theoretical Economics* 17.2 (2022): 561-585.

[46] Kolotilin, Anton, Roberto Corrao, and Alexander Wolitzky. "Persuasion and Matching: Optimal Productive Transport." *Journal of Political Economy* (2024): Forthcoming.

[47] Klein, Timo. "Autonomous Algorithmic Collusion: Q-Learning under Sequential Pricing." *RAND Journal of Economics* 52.3 (2021): 538-558.

[48] Kuvalekar, Aditya, Elliot Lipnowski, and Joao Ramos. "Goodwill in Communication." *Journal of Economic Theory* 203 (2022): 105467.

[49] Lamba, Rohit, and Sergey Zhuk. "Pricing with Algorithms." *arXiv preprint* arXiv:2205.04661 (2024).

[50] Marshall, Robert C., and Leslie M. Marx. *The Economics of Collusion: Cartels and Bidding Rings*. Mit Press, 2012.

[51] Martin, Simon, and Alexander Rasch. "Demand Forecasting, Signal Precision, and Collusion with Hidden Actions." *International Journal of Industrial Organization* 92 (2024): 103036.

[52] 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.

[53] McAfee, R. Preston, and John McMillan. "Bidding Rings." *American Economic Review* (1992): 579-599.

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

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

[56] 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.

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

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

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

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

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

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

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

[64] Smolin, Alex and Takuro Yamashita. "Information Design in Smooth Games." *Working Paper*, Toulouse School of Economics (2023).

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

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

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

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

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

[70] Vives, Xavier. *Oligopoly Pricing: Old Ideas and New Tools.* MIT Press, Cambridge, Mass., 2001.